

Wii U CPU Profiler Manual

Wii U CPU Profiler

2015/07/02

Version 1.00

**The content of this document is highly
confidential and should be handled accordingly.**

CONFIDENTIAL

These coded instructions, statements, and computer programs contain proprietary information of Nintendo and/or its licensed developers and are protected by national and international copyright laws. They may not be disclosed to third parties or copied or duplicated in any form, in whole or in part, without the prior written consent of Nintendo.

Contents

1	Overview.....	11
1.1	Installation.....	11
1.2	Requirements.....	11
1.3	Limitations.....	11
1.4	Resources.....	12
2	Taking a Profile.....	13
2.1	Setup.....	13
2.2	Profiling.....	14
3	User Logged Data and Events.....	15
3.1	Heartbeats (Frames).....	15
3.1.1	Frame Marking on Wii U.....	15
3.1.2	Inferred Heartbeats.....	15
3.2	User Logged Data.....	16
3.2.1	User Logged Data on Wii U.....	16
3.3	Marked Code Blocks.....	16
4	Profiler Workflows.....	17
4.1	Quickly Understanding Code Flow.....	17
4.2	Finding Spiky Functions.....	17
4.3	Identifying the Cause of a Slow Frame.....	18
4.4	Find all Sampled Functions at a Particular Time.....	18
4.5	Measure the Time of a Sample Run in the Sample Graph.....	18
5	Ribbon Bar.....	20
5.1	Home Tab.....	20
5.1.1	Connection.....	20
5.1.2	File.....	20
5.1.3	Help.....	20
5.1.4	Units.....	20
5.1.5	Color Coding.....	21
5.2	Analysis Tab.....	22
5.2.1	Story Analysis.....	22
5.2.2	Spike Detection.....	23
5.2.3	Quick Filters.....	24
5.3	Heartbeats Tab.....	26
5.3.1	Group Selection.....	26
5.3.2	Heartbeat Cores.....	27
5.4	Sampled Profile Tab.....	28
5.4.1	Profile Buttons.....	28
5.4.2	Sampling Strategy and Rate Buttons.....	28
5.4.3	Data per Sample.....	29
5.4.4	Cores.....	30
5.4.5	Recording Marked Code Blocks.....	30

5.5	Instrumented Profile Tab.....	30
5.5.1	Allowing for Instrumentation.....	31
5.5.2	Profile Buttons.....	31
5.5.3	Choose a Function to Instrument.....	31
5.5.4	Choose a Performance Counter Group.....	32
5.5.5	Recording Marked Code Blocks.....	32
5.6	Directories Tab.....	32
5.6.1	SLC Directory.....	32
5.6.2	Button Directory.....	33
5.6.3	Code Directory.....	33
5.7	Quick Ribbon Bar.....	33
6	Functions Tab.....	34
6.1	Sorting Functions.....	34
6.2	Functions Tab Toolbar.....	34
6.2.1	Core Filtering and Column Heading.....	34
6.2.2	Deselect All Button.....	34
6.2.3	Select Top Button.....	34
6.2.4	Diff Total and Diff Self Column Headings.....	34
6.2.5	Threading Filtering and Thread Column Heading.....	35
6.2.6	Module Column Heading.....	35
6.2.7	Opcodes Column Heading.....	35
6.2.8	Args Button.....	35
6.2.9	Copy Button.....	35
6.2.10	System Classification Filter.....	36
6.2.11	Selected Only Button.....	36
6.3	Filtering Functions.....	36
6.3.1	Quick Regular Expression Primer.....	37
6.3.2	Saved Filters.....	37
6.4	Selecting Functions.....	37
6.4.1	Alternative Selection Methods.....	37
6.5	Right-Click Context Menu.....	38
6.5.1	Copy Function Name to Clipboard.....	38
6.5.2	Deselect All Other Functions.....	38
6.5.3	Show Assembly.....	38
6.5.4	Show Details.....	38
6.6	Units.....	38
6.7	Resizing the Window.....	38
6.8	Reading Obscured Function Names.....	38
6.9	Multiple Function Names on a Single Line.....	38
6.10	RPL Internal Functions.....	39
6.11	Unresolved Functions.....	39
7	Last Tab.....	40
8	Threads Tab.....	41
8.1	Sorting Threads.....	41
8.2	Threads Tab Toolbar.....	41
8.2.1	Core Filtering and Column Heading.....	41
8.2.2	Deselect All Button.....	41
8.2.3	Top Button.....	41
8.2.4	Copy Button.....	41
8.3	Selecting Threads.....	41

8.3.1	Alternative Selection Methods.....	42
8.4	Filtering Threads.....	42
8.4.1	Saved Filters.....	42
8.5	Resizing the Window.....	42
9	Instrumented Tab.....	43
9.1	Instrumented Tab Toolbar.....	43
9.1.1	Deselect All Button.....	43
9.1.2	Expand All Button.....	43
9.1.3	Collapse All Button.....	43
9.1.4	Select Similar Button.....	43
9.1.5	Copy Button.....	43
9.2	Instrumented Items.....	43
9.3	Selecting Instrumented Items to Graph.....	44
9.3.1	Alternative Selection Methods.....	44
9.4	Resizing the Window.....	44
10	Counters Tab.....	45
10.1	Types of Counters.....	45
10.1.1	Performance Counters.....	45
10.1.2	OS Statistics.....	45
10.1.3	Usage Statistics.....	45
10.1.4	GPU Statistics.....	45
10.1.5	User Logs.....	48
10.2	Sorting Counters.....	48
10.3	Counters Tab Toolbar.....	48
10.3.1	Core Filtering and Column Heading.....	48
10.3.2	Group Filtering and Column Heading.....	48
10.3.3	Deselect All Button.....	49
10.3.4	Select Similar Button.....	49
10.3.5	Copy Button.....	49
10.4	Filtering Counters.....	49
10.4.1	Saved Filters.....	49
10.5	Selecting Counters.....	49
10.5.1	Alternative Selection Methods.....	50
10.6	Resizing the Window.....	50
10.7	Units.....	50
11	Info Tab.....	51
11.1	Info Tab Toolbar.....	51
11.1.1	Expand All Button.....	51
11.1.2	Collapse All Button.....	51
11.1.3	Copy Button.....	51
11.2	Environment.....	51
11.3	Sampling Settings.....	51
11.4	Profile Statistics.....	52
11.5	Module Statistics.....	52
11.6	Inferred Heartbeats.....	52
11.7	Application Information.....	53
11.8	Profiler Information.....	53
11.9	Notes.....	53

12	Checkers Tab.....	54
12.1	Checkers Tab Toolbar.....	54
12.1.1	Manage Checkers.....	54
12.1.2	Status.....	54
12.2	Right-Click Context Menu.....	54
12.2.1	Select Function.....	54
12.2.2	Show Details.....	55
12.2.3	Show Assembly.....	55
12.2.4	Load per Frame Analysis.....	55
12.3	Severity.....	55
12.4	Checker Results.....	55
12.4.1	Rerun Checker Button.....	55
12.4.2	Copy Text Button.....	55
12.4.3	Copy Table Button.....	55
12.4.4	Disable Checker Button.....	55
12.4.5	Summary.....	56
12.4.6	Configurable Parameters.....	56
12.4.7	Reset to Defaults Button.....	56
12.4.8	Results Table.....	56
12.4.9	Explanation.....	56
12.4.10	Suggestions.....	56
13	Call Tree Tab.....	57
13.1	Call Tree Tab Toolbar.....	57
13.1.1	Prefix Selection.....	57
13.1.4	Auto Expanding.....	57
13.1.5	Copy Button.....	57
13.1.6	Copy All Button.....	58
13.2	Finding Functions.....	58
13.2.1	Saved Filters.....	58
13.3	Function Selection.....	58
13.4	Right-Click Context Menu.....	59
13.4.1	Manual Expanding.....	59
13.4.2	Show Assembly.....	59
13.4.3	Show Details.....	59
13.5	Partial Call Tree.....	59
13.6	Units.....	60
14	Sample Graph Tab.....	61
14.1	Sample Graph Tab Toolbar.....	61
14.1.1	Mouse Mode.....	61
14.1.2	Deselect All Button.....	62
14.1.3	Icicle Graph.....	62
14.1.4	Icicle Selected Graph.....	63
14.1.5	Stack Graph.....	63
14.1.6	Load Graph.....	63
14.1.7	Heatmap Mode.....	63
14.1.8	Self Samples / Total Samples / Mixed Samples.....	63
14.1.9	Adjusting Name Widths.....	64
14.1.10	Combined Graphs.....	64
14.1.11	Graph Tool Tips.....	64
14.1.12	Line Highlight.....	64

14.1.13	Sample Gaps.....	65
14.1.14	Samples, Load, and Counters Scale.....	65
14.1.15	View Control.....	65
14.2	Timeline.....	65
14.3	Zoom Control.....	65
14.4	Right-Click Context Menu.....	66
14.4.1	Expand Direct Parents.....	66
14.4.2	Expand Direct Children.....	66
14.4.3	Deselect Direct Children.....	66
14.4.4	Reorder.....	66
14.4.5	Deselect.....	66
14.4.6	Move.....	66
14.4.7	Show Assembly.....	67
14.4.8	Show Details.....	67
14.4.9	This Core At This Time.....	67
14.4.10	All Cores At This Time.....	67
14.4.11	Select Only This Heartbeat Frame.....	67
15	Code Coverage Tab.....	68
15.1	Functions Seen During Profiling.....	68
15.1.1	Copy Button.....	68
15.2	Functions Not Seen During Profiling.....	68
15.2.1	Copy Button.....	68
15.3	Filtering Functions.....	68
15.3.1	Saved Filters.....	69
15.4	Right-Click Context Menu.....	69
15.4.1	Copy Function Name to Clipboard.....	69
15.4.2	Select Function.....	69
15.4.3	Show Assembly.....	69
15.4.4	Show Details.....	69
16	Assembly Tab.....	70
16.1	Assembly Tab Toolbar.....	70
16.1.1	Copy Button.....	70
16.1.2	View Controls.....	70
16.1.3	Show Source.....	70
16.2	Selecting a Function to View.....	71
16.3	Assembly View.....	71
16.3.1	Highlighting.....	71
16.3.2	Sorting.....	71
16.4	Right-Click Context Menu.....	72
16.4.1	Select Target Function.....	72
16.4.2	Jump to Target.....	72
16.4.3	Show Target Details.....	72
16.4.4	Copy Target Function Name.....	72
16.4.5	View Data As.....	72
17	Console Tab.....	73
17.1	Console Tab Toolbar.....	73
17.1.1	Enable Button.....	73
17.1.2	Clear Button.....	73
17.1.3	Copy Button.....	73
17.1.4	Text Format Selection.....	73

18	Function Details Window.....	74
18.1	Selecting a Function to Display.....	74
18.2	Function Details Operations.....	74
18.2.1	Show Assembly.....	74
18.2.2	Select Function.....	74
18.2.3	Traversing Function Calls.....	75
19	System Classifications.....	76
19.1	Customizing Classifications.....	76
20	Performance Counter Groups.....	77
20.1	Performance Counters Disabled.....	77
20.2	Per-Core Performance Counter Groups.....	77
20.2.1	Cycles and Instructions.....	77
20.2.2	Branch Prediction Performance.....	77
20.2.3	Why Branch Prediction Failed.....	78
20.2.4	Cache and Memory Performance.....	79
20.2.5	Is Data Access the Bottleneck?.....	79
20.2.6	L1 Cache Performance.....	79
20.2.7	L2 Cache Performance.....	80
20.2.8	Cache Misses to Memory.....	80
20.2.9	Outbound Cache Writes.....	81
20.2.10	High Cost Cache Writes.....	81
20.2.11	Cache Line Push Counts.....	81
20.2.12	L1 Sharing Relationships.....	82
20.2.13	L2 Sharing Relationships.....	82
20.2.14	TLB Statistics.....	82
20.3	Chip-Wide Performance Counter Groups.....	83
20.3.1	Inter-Core Paradox.....	83
20.3.2	Cache Sharing on Chip.....	83
20.3.3	Chip Data Transfers.....	83
20.4	Correlation Between Sampled Functions and Performance Counters.....	84
20.5	Performance Counter Descriptions.....	84
21	Command Line Options.....	87
21.1	Command Line Format.....	87
21.2	Available Options.....	88
22	Scripting.....	89
22.1	Windows PowerShell.....	89
22.2	Cmdlets.....	89
22.2.1	PROFILER-ANALYZE.....	89
22.2.2	PROFILER-ASSEMBLY.....	90
22.2.3	PROFILER-CALLSTACKS.....	90
22.2.4	PROFILER-COPY.....	90
22.2.5	PROFILER-FILTER.....	91
22.2.6	PROFILER-INSTRUMENT.....	92
22.2.7	PROFILER-MARKED.....	93
22.2.8	PROFILER-OPEN.....	93
22.2.9	PROFILER-PERFCOUNTERS.....	93

22.2.10	PROFILER-SAMPLE.....	93
22.2.11	PROFILER-SAMPLEGRAPH.....	94
22.2.12	PROFILER-SAVE.....	95
22.2.13	PROFILER-SELECT.....	95
22.2.14	PROFILER-SHOWTAB.....	96
22.2.15	PROFILER-START.....	96
22.2.16	PROFILER-STOP.....	96
22.2.17	PROFILER-SYNC.....	97
22.2.18	PROFILER-TREECONTROL.....	97
22.2.19	PROFILER-UNITS.....	97
22.2.20	PROFILER-UNSYNC.....	98
22.2.21	PROFILER-WAIT.....	98
22.2.22	PROFILER-WAITFOR.....	98
22.2.23	PROFILER-CAFE-CORES.....	99
22.2.24	PROFILER-CAFE-GPUSTATS.....	99
22.2.25	PROFILER-CAFE-OSSTATS.....	99
22.2.26	PROFILER-CAFE-SETDIR.....	99
22.3	Argument Values.....	100
22.3.1	Copy Extensions.....	100
22.3.2	Profiler Wait Events.....	100
22.3.3	Rate.....	100
22.3.4	Scripting Directory Type.....	100
22.3.5	Scripting Perf Counter.....	100
22.3.6	Scripting Profile Type.....	101
22.3.7	Scripting Sample Strategy.....	101
22.3.8	Scripting Tabs.....	101
22.3.9	Scripting Tree Control Actions.....	102
22.3.10	Scripting Unit Type.....	102
22.3.11	Scripting.SampleGraph.CombinedLineDisplay.....	102
22.3.12	Scripting.SampleGraph.NameDisplay.....	102
22.3.13	Scripting.SampleGraph.SampleDisplay.....	102
22.3.14	Spike Detection Types.....	102
22.3.15	Story Types.....	102
22.4	Other Useful Commands.....	103

23 Wii U CPU Profiler Game API..... 104

24 Assembly Files and Functions..... 105

25 Troubleshooting..... 106

25.1	Profiler crashes.....	106
25.2	Sample Graph draws very slow and has overlapping frame marker.....	106
25.3	Sample Graph does not appear.....	106
25.4	Sample Graph appears to be a frame behind.....	106
25.5	Sync button stays grayed out.....	106
25.6	Security warning pop-up after starting profiler.....	106
25.7	Inability to take an Instrumented Profile.....	107
25.8	Unresolved functions.....	107
25.9	Unable to locate the compiler directory.....	107
25.10	C++ functions are not demangled.....	107

26 Known Issues..... 108

26.1	Call Tree cores disappear after Collapse All.....	108
26.2	Call Tree seems wrong.....	108
26.3	Occasionally wrong SDK function reported with sampling.....	108
27	Glossary.....	109
28	Revision History.....	111

1 Overview

The Wii U CPU Profiler is a statistical sampling profiler that helps detect which functions are using significant CPU resources. The profiler also has the ability to instrument one function at a time to accurately and precisely measure its use of CPU resources.

1.1 Installation

Unzip the profiler package to a directory of your choice.

Wii U Specific Installation Instructions

Open the `runtime\Cafe\cafe_sdk` folder and copy the contents into your SDK install directory so that they can be referenced during compilation by your game.

When running the profiler for the first time, be sure to set up the **Directories** as this will allow the profiler to find the SDK RPLs, thereby removing most UNRESOLVED functions. See [SLC Directory](#) for more information.

1.2 Requirements

The profiler requires the following:

General Requirements

- Windows 7 64-bit
- DirectX 10
 - For all of the required libraries, the DirectX End-User Runtimes (June 2010) must be separately installed onto the computer. The DirectX installer can be found for download from <http://www.microsoft.com/download/en/details.aspx?id=8109>.
- .NET 4.0 Runtime (search Microsoft.com for ".net 4.0 runtime" or visit <http://www.microsoft.com/en-us/download/details.aspx?id=17851>).

Wii U Specific Requirements

- CAFE SDK 2.11.X or 2.12.X.
 - The profiler is tied to specific major releases of the SDK. Please ensure that you are using the correct profiler version with the correct SDK version(s).
- The compiler makes use of certain compiler features for the **Assembly** tab and for demangling C++ function names. The profiler attempts to locate the directory for the compiler to determine where to find these tools by searching for the environment variable `GHS_ROOT`. If that fails, it looks for the install directory set up by the installer.

1.3 Limitations

The profiler has a few limitations:

- Custom assembler files must have functions delineated with the `.type` and `.size` directives. See [Assembly Files and Functions](#) for more information.
- The SDK defaults to having code protection enabled. To take an instrumented profile, debugging mode must be enabled. See [Allowing for Instrumentation](#) for more information on how to enable this mode.

- If a thread is created after a profile has started, the profiler cannot automatically get its name. In this case `PROFILERSetThreadName` must be called. A maximum number of 255 thread names can be registered in the profiler at once. This includes threads that were running when profiling started. The stored names are reset at the start of each profile.
- RPL internal function names are based on location in the RPL. As such, it may not be possible to effectively compare internal functions between different RPL builds. See [RPL Internal Functions](#) for more information.

1.4 Resources

The profiler requires a large memory buffer of size `PROFILER_MINBUFFERSIZE` (defined in `profiler.h`, currently set at 64MB), which you must pre-allocate and pass into `PROFILERInitialize`. The memory is used to store sample data from each core, plus some additional bookkeeping information. If you wish to take longer profiles (collect more samples per profile), pass a larger buffer into `PROFILERInitialize`.

For its operation, the profiler creates a total of 5 threads. Those threads are detailed as followed:

Table 1-1: Profiler Thread Usage

Thread Name	Core	Priority
Main Thread	Any	1
Core 0 Sampler	0	2
Core 1 Sampler	1	2
Core 2 Sampler	2	2
Comm Thread	2	2



Note: Thread names are prefixed by the text "Wii U CPU Profiler - ".

A more detailed list of resource usage will be made in a future revision of the profiler.

2 Taking a Profile

2.1 Setup

To initially setup the profiler for use with your game, please follow these steps.

1. Add the profiler library to the project.

If using the Makefile system used by the SDK, this can be accomplished with the following:

```
MODULE_LIBNAMES += profiler.a
```

2. Include the profiler header file in your source code. Add: **#include <cafe/profiler.h>**.

```
#include <cafe/profiler.h>
```

3. Initialize the profiler in your code.

```
void *buffer = MEMAllocFromDefaultHeap(PROFILER_MINBUFFERSIZE);
BOOL profInitResult = PROFILERInitialize(buffer, PROFILER_MINBUFFERSIZE);
```

The first argument is a memory buffer and the second argument is the size of the buffer. The size of the buffer to allocate must be of at least `PROFILER_MINBUFFERSIZE` bytes.

4. Add a frame reference marker to each core (generically called a "heartbeat" in the profiler).

For example, on core 1 you might put the following in your main loop:

```
PROFILERHeartbeat(PROFILEREventHeartbeatName_Main,
                  PROFILEREventHeartbeatType_Beat);
```

The marker `PROFILEREventHeartbeatName_Main` is special as it is used to determine the frame rate of your game (only use it on one core in a single place). More information on instrumenting heartbeats can be found in the API man pages.

5. Start the Game or Application.

It is strongly recommended that you use the following command as these options are required for **Instrumented** profiling.

```
caferun -d multi -R RpxPath
```

For more information, please see [Allowing for Instrumentation](#).

6. Start the Profiler by running the application "Nintendo CPU Profiler.exe" located in the bin folder.
7. Sync the Profiler by clicking the blue **Sync** button on the ribbon bar.
 - Find the files requested by any dialog boxes that appear.
8. Set the **RPL Directory**.

In the profiler, click on the **Directories** tab. Find the box marked **RPL Dir** and set it to the SDK's RPL directory. For more information please see [SLC Directory](#).

2.2 Profiling

1. Start the Game or Application.

It is strongly recommended that you use the following command as these options are required for **Instrumented** profiling.

```
caferun -d multi -R RpxPath
```

For more information, please see [Allowing for Instrumentation](#).

2. Start the Profiler by running the application "Nintendo CPU Profiler.exe" located in the `bin` folder.
3. Sync the Profiler by clicking the blue **Sync** button on the ribbon bar.
 - Find the files requested by any dialog boxes that appear.
4. On the **Sampled Profile** tab, select the sampling strategy, rate, and other options.
5. Click on the green **Start** button to start profiling.
6. Stop Profiling
 - To manually stop, click on the red **Stop** button.
 - To automatically stop after a set amount of time, use the drop-down just to the right of **Start** and **Stop** buttons to select a time.
 - Alternatively, profiling will automatically stop once the recording buffer is full (this may take awhile).
7. Examine the profile in the **Functions**, **Info**, **Call Tree** (if **Callstack** mode was enabled), and **Code Coverage** tabs.
8. Optional: Click the **Save** button to save the profile.

3 User Logged Data and Events

While the profiler can be used with very little manual integration, it can be extremely useful to augment a profile with user logged data and events. The following describes available options.

3.1 Heartbeats (Frames)

When looking at data over long time periods, it is helpful to subdivide the data into meaningful groups, such as frames. Since multiple systems and different cores may operate in different rhythms, we've come up with the broader term "heartbeat" to denote any regularly occurring rhythm or framing that is useful to track.

If a heartbeat was logged during a profile, it will appear in one of the drop down boxes in the **Heartbeats** tab. When selected, the heartbeat will be graphed for that particular core. You'll notice that there are always two other heartbeats in the drop down boxes: **Fixed 60Hz Intervals** and **Fixed 1ms Intervals**. These heartbeats weren't actually recorded in the profile data, but you can select them to overlay timing information onto a particular core. Please note that these two heartbeats are unlikely to correlate with your data and are only offered as a timing reference.

If you wish to focus on particular heartbeat frames, you can select them either by clicking and dragging over the frames you want to select or by clicking on a particular frame rate in the **Heartbeats** tab (for example, the **60Hz** toggle button).

3.1.1 Frame Marking on Wii U

The profiler API allows you to log multiple heartbeats on each core. The **VSynchHeartbeat** is logged automatically, but all others must be logged by the game. Up to 9 different user heartbeats can be logged per core, as listed in the enumeration `PROFILEREventHeartbeatName` in `profiler.h`.

The most important heartbeat to log is `PROFILEREventHeartbeatName_Main`. This heartbeat is used by the profiler to draw a frame rate graph at the bottom of the **Sample Graph**. However, it is extremely important to only log the main heartbeat on your rendering core and not on any other cores (as this confuses the frame rate graphing in the GUI).

With each heartbeat event, you must also chose the type.

- `PROFILEREventHeartbeatType_Beat`
- `PROFILEREventHeartbeatType_BeginWait`
- `PROFILEREventHeartbeatType_EndWait`

Of these, `PROFILEREventHeartbeatType_Beat` is typically the most useful.

While all three versions will be logged and can be graphed individually, we have not yet reached our intended vision with the graphing (to have all three versions graphed at once).

3.1.2 Inferred Heartbeats

While it is still critical for you to log the Main heartbeat, the profiler will infer heartbeats on each core based on function behavior. If there are functions that regularly execute on a periodic basis, the profiler will use the most periodic functions as an indicator of a heartbeat rhythm on that core. Inferred heartbeats are more likely to be detected at higher sampling rates (i.e. 100x or higher), since there is more data per frame to analyze. These inferred heartbeats are then selectable in the **Heartbeats** tab. Additionally, the **Info** tab has a list of all inferred heartbeats that were detected. Unfortunately, since these heartbeats are educated guesses based on the data, they are unlikely to actually line up with the conceptual beginning of a frame. If this is a problem, it is recommended that you explicitly log the heartbeat to accurately mark where the frame begins.

3.2 User Logged Data

Sometimes it can be very powerful to see game data graphed against profiler data. For this case, the profiler offers the ability to log custom data.

Data that is logged will be displayed on the **Counters** tab. When selected on the **Counters** tab, it will be graphed in the **Sample Graph**.

Please see the API documentation for the exact details on how to log data.

3.2.1 User Logged Data on Wii U

- **Pulse**

Log a moment in time that marks an interesting event.

- **Unsigned Int (32-bit)**

Log a positive integer that represents something in the game.

- **Float (32-bit)**

Log a floating-point number that represents something in the game.

Supplied with the profiler package is a sample demo ("benchmarkProfiled") that shows how to log graphics benchmarks during a profile.

3.3 Marked Code Blocks

Many game developers wrap their game systems with timers so that exact running times can be easily tracked during gameplay. The profiler supports logging these marked blocks of code with the added benefit of being able to track CPU performance counters along with the execution time. When coupled with the profiler, you can use the profiler to graph all of these metrics along with a profile. Please see the [Wii U CPU Profiler Game API](#) documentation for details on integrating the code block feature.

Once your code has been manually instrumented with code blocks, you must enable the recording of them in the profiler. To do so, use the **Marked** button on either the **Sampled Profile** tab or the **Instrumented Profile** tab. You can then select which performance counter group to record along with your marked code blocks.

Once recorded in a profile, the marked code blocks will be listed on the **Instrumented** tab. When selected on the **Instrumented** tab, they will be graphed in the **Sample Graph**.

4 Profiler Workflows

When using the profiler, there are many possible workflows to discover or track down problems. The following is a list of interesting use cases.

4.1 Quickly Understanding Code Flow

Use the **Story Analysis** buttons on the **Analysis** ribbon bar tab to quickly get a handle on the code flow within a frame. This is the fastest way to understand the behavior within a frame.

Setup

- Record at least **1000x per frame**.
- Record a **Callstack** profile.
- A long profile is not necessary since you will be looking at a particular frame.

Procedure

1. In the **Functions** tab, select the core you want to examine using the Core Filter (described in [Core Filtering and Column Heading](#)).
2. In the **Analysis** tab in the ribbon bar, click on the **Brief** button in the **Story Analysis** group.
3. In the **Sample Graph** tab, zoom into a typical frame or a frame that you want to deeply understand. At this point you can see the behavior from the start of the frame to the end.



Tip: There are a few things that can help make it easier to view the data.

- Click the **Load** button in the **Sample Graph** toolbar to temporarily hide the **Load per Frame** graph. Click it again to bring it back.
 - Use the **Samples** drop down in the **Sample Graph** toolbar to adjust the vertical spacing of samples. A larger vertical spacing makes it easier to visually track a sample back to its function name of the left.
 - You can always order the samples by chronological order by right clicking on the **Sample Graph** and choosing **Reorder > Reorder Chronologically**.
 - You can quickly drill down into a function's children by double clicking on the function line in **Sample Graph**. To undo what you just expanded, remain on the same function line, right click, and choose **Deselect Direct Children**.
4. Optional: Select the frame you are looking at and click **Brief** again. This will only show interesting functions on this frame, sorted correctly for this frame. To select a frame, change the drop down mouse mode on the far left of the **Sample Graph** toolbar to **Select Frames** and click on your desired frame once.
 5. To get a more detailed story, click on **Full**, **Fine**, or **Top** in the **Story Analysis** group.
You can adjust the settings on the left side of the **Story Analysis** group or you can adjust the parameters within a particular **Story Analysis** button by clicking on the name to expand the options. See [Story Analysis](#) for more information on how to do this.

4.2 Finding Spiky Functions

With one click you can find functions that spike occasionally.

Setup

- Record at least **1000x per frame**.
- Record a **Callstack** profile.

Procedure

Click the **Rare** button within the **Spike Detection** group on the **Analysis** ribbon bar tab.



Tip: Click on the **Rare** label to expand the options and increase the **Max Spikes**. Then click on the **Rare** button again to reanalyze the profile.

4.3 Identifying the Cause of a Slow Frame

With enough samples per frame, the exact cause of a slow frame can be determined with one button click.

Setup

- Record at least **1000x per frame**.
- Record a **Callstack** profile.
- Record a profile long enough to have a slow frame.

Procedure

1. Optional: In the **Functions** tab, select the core you want to examine before clicking **Bad Frame**.
2. Click the **Bad Frame** button within the **Spike Detection** group on the **Analysis** ribbon bar tab.



Tip: Additional options are available by expanding the **Bad Frame** options.

- Increase the number of bad frames to detect.
- Enable **Use Selected Frames**, select the frames you are concerned about (see [Mouse Mode](#)), and click **Bad Frame**. Only functions that perform worse on those frames compared with all other frames will be selected.

4.4 Find all Sampled Functions at a Particular Time

You might want to know what functions were sampled at a particular time. For example, maybe a CPU performance counter indicates poor IPC at a particular time.

Setup

- Record a **Callstack** profile.
- A long profile is not necessary since you will be looking at a particular moment in time.

Procedure

1. Optional: In the **Functions** tab, deselect all functions. This will help reduce clutter in the **Sample Graph**.
2. In the **Sample Graph** tab, find a time that is interesting.
3. Right click on the **Sample Graph** and choose **This Core at This Time > Find All Sampled Functions**.

4.5 Measure the Time of a Sample Run in the Sample Graph

To measure a sample run in the **Sample Graph**, change the mouse mode to **Measure Time** in the **Sample Graph** toolbar (far left side).

Setup

- Change the mouse mode to **Measure Time** in the **Sample Graph** toolbar (see [Mouse Mode](#)).

Procedure

In the **Sample Graph**, left click at the start of the sample and drag the mouse to the end of the sample.



Remember: The mouse behavior changes in different mouse modes, with different shortcuts available to accomplish tasks.

- In the **Measure Time** mouse mode, you can scroll the window with the mouse by holding the Ctrl key while left clicking and dragging.
- In the **Scroll Graph** mouse mode, you can measure time if you hold the Ctrl key while left clicking and dragging.
- To make your selection stay and select multiple areas, use the **Select Time** mouse mode. Be aware that this will reevaluate the numbers in the **Functions** and **Call Tree** tabs based on the selected regions.

The measured time will be shown at the top of the selected area (if **Load per Frame** is hidden, then the time will be shown at the bottom of the selected area). When you release the left mouse button, the measured area will disappear.

5 Ribbon Bar

This section details the various controls and settings available in the ribbon bar menu.

To hide the ribbon, click the up arrow at the top right. Once hidden, the ribbon bar will show only when the mouse cursor rolls over the ribbon bar tabs. To expand the ribbon so that it stays shown, click the down arrow at the top right.

5.1 Home Tab

The **Home** tab contains the primary profiler operations for connecting to the dev kit and manipulating the profiles.

5.1.1 Connection

In order to connect to a dev kit, you must first press the **Sync** button. Pressing this button starts the following three step synchronization process:

1. Selection of a dev kit. If only one dev kit is available that kit will be connected to automatically. If more than one kit is available a window will appear with a drop down with the list of available dev kits.
2. You will be prompted to locate the executable that you wish to profile. This is used to cross-reference function addresses with function names.
3. The Dev Kit is pinged and the allocated buffer size is sent to the PC.

Once the sync process has completed, the **Start** profile button and the **Unsync** button will become active.

Press the **Unsync** button if you want to profile a different executable or if you want to connect to a different dev kit.

Current connection information is shown in the **Status Bar** at the bottom of the application.

5.1.2 File

The **Open** button allows you to open a previously saved profile.

The **Save** button saves the current profile to disc.

The **Script** button loads and runs scripts that can automate profiling.

5.1.3 Help

The **Manual** button will launch this document (Wii U CPU Profiler Manual). The profiler expects the manual to be one directory up from the profiler executable. In the case that the profiler can't locate the manual, the error message will display the manual filename and directory that it expects.

5.1.4 Units

These buttons control the units used to represent the amount of time spent in each function (for sampled profiles only). The selected units will be reflected in any tab that shows profile data, such as the **Functions** tab, the **Call Tree** tab, or in the tool tips of the **Sample Graph** tab. The various units allow several ways to analyze and compare functions.

Percent

This is the percent of time the function ran compared with all other functions during the selected time period. For example, if a function was sampled 100 times out of 400 total samples, then 25% of the

execution time was spent in this function. **Percent** is useful when comparing relative function times. Enable the **Error** button in order to see the margin of error on the values.

Time

This is an approximation of the time spent in each function (the approximation is calculated with statistics and is not directly measured using a timer). If one or more frames are selected or a period of time is selected, then the time is an estimate over that selection. If frames or times are not selected, then the time is an estimate over the entire duration of the profile.

Use **Time** to understand how long a function actually executed, which can be useful if you want to stay within a particular time budget or frame rate. For example, you can use this feature to determine if a particular module or system is going over its budgeted time (such as 2ms each frame for the AI to update). Please be aware that **Time** will be more accurate if you use a very low sampling rate over a very long period of time. As profiler overhead increases, **Time** will become less accurate, so a low sampling rate is critical if you want to use this metric. Also, be aware that selecting individual frames will decrease the precision/accuracy as the margin of error will increase. Enable the **Error** button to see the margin of error (note that this error calculation does not take into account profiler overhead, just statistical variation).

Avg

This is an approximation of the time spent in each function divided by the number of frames. If no frames exist for a given core, then this is the same as the **Time** unit (nothing is averaged). If the core has frames, but none are selected, then time is divided by the total number of recorded frames on that core. If the core has selected frames, then time represents the time on those selected frames divided by the number of currently selected frames. Enable the **Error** button to see the margin of error.



Note: With **Avg** selected, the sorting in the **Total** and **Self** columns in the **Functions** tab may appear wrong. This happens because the functions on different cores are divided by a different number of frames, yet the sorting is based on importance (absolute time). However, sorting is correct for all functions on a given core (sorting is correct if filtered by core).

Samples

This provides the actual number of samples taken in a function. This is the basic unit that the above units are calculated from. Enable the **Error** button to see the margin of error.

Error

This button is independent of the other **Units** buttons. It controls whether or not the margin of error is displayed next to the **Percent**, **Time**, **Avg**, or **Samples** values. The margin of error is useful to determine the precision of the results. For example, if a single frame is selected, the margin of error tends to be very large. As a result, the **Percent** or **Time** values may look precise (for example, 8.2%), but the margin of error can tell you the exact precision (for example, $\pm 5.8\%$), which is important when comparing values or assessing performance. The **Error** button's state is saved in the profiler's settings so that it will remember the last selected state when you relaunch the profiler.

5.1.5 Color Coding

The **Color Coding** group controls the color scheme for selected functions within the profiler.

Function

Function assigns every selected function it's own color (with colors being recycled after every 15 selected functions).

System

System is the other choice, which gives each selected function a color based on which classification system it belongs to. Functions are automatically classified as belonging to one of over a dozen different systems, such as Physics, Graphics, or Audio. For more information on the classification system, please see [System Classifications](#).

If you want selected functions to match their respective color within the **Iceberg Graph**, then you should select **System** for the **Color Coding** setting.

If **System** is chosen, the key that explains each color classification is displayed above each core heading in the **Sample Graph**.

5.2 Analysis Tab

The **Analysis** tab has three groups of tools to help you quickly analyze a profiler: **Story Analysis**, **Spike Detection**, and **Quick Filters**.

5.2.1 Story Analysis

The **Story Analysis** tools help select and order functions to tell an abstract story about what happens over the course of a frame.

To use these tools, you must have taken a **Callstack** profile, you should zoom into a single frame within the **Sample Graph**, and it helps to only look at a particular core (select a single core in the **Functions** tab toolbar).

Story Analysis Options

These options help refine the selected functions.

Exclude Parents: By default, the parents of interesting functions are also selected. This toggle will exclude these parents from the results. This option has no effect on **Top Story Analysis**.

Exclude Irregular: By default, all functions can be included as part of the story. This toggle will exclude functions that do not execute on a regular basis.

Match Heartbeat: By default, all functions can be included as part of the story. This toggle will exclude functions that do not appear at an interval similar to the frame rate.

5.2.1.1 Brief Button

This button selects functions that tell a brief story about what happens over the course of a frame. It is more effective if you first select a single core to look at (see [Core Filtering and Column Heading](#)) before hitting this button. Zoom into a single frame in the **Sample Graph** to see the story.

Brief Button Options

The settings of **Max** and **Min** in the button options control which functions are selected. If you change these settings, they will not be saved permanently. However, there are three preset setting groups that you can try: **Short**, **Long**, and **Detailed**. These presets offer varying levels of detail in the generated story, going from most high-level (selects minimal number of key functions) to extremely detailed (selects many functions to show extreme, yet significant, detail).

5.2.1.2 Top Button

This button is different from the previous three (**Brief**, **Full**, and **Fine**) and implements a slightly different algorithm. This button will select the top functions, exclude particular ones based on settings, then order them chronologically.

Top Button Options

The settings in the **Top** button options will help control what gets selected. The **Max Results** setting is very useful for controlling the number of functions that get selected (however the final results may be lower than this number based on other options that exclude functions). If the results contain many similar functions, you can eliminate similar ones with either the **Exclude Similar Children** or **Exclude Similar Parents** toggle buttons.

5.2.2 Spike Detection

The **Spike Detection** tools help quickly find potentially troublesome functions that you should further investigate.

Spike Detection Options

These options help refine the selected functions.

Max Results: This determines how many functions are selected when **Rare**, **Burst**, **Flutter**, or **Bad Frame** are clicked. One strategy is to make this number very low (such as 5 or 10) to hyper focus on the worst spiking functions. Another strategy is to make this number very big (such as 100 or 200) and see if anything interesting shows up.

Exclude Similar Children: If this is enabled, children functions that are too similar to their parent will be discarded from the results. This helps to understand which high-level systems are causing the spikes, since parent functions tend to be named after the high-level systems they represent.

Exclude Similar Parents: If this is enabled, parent functions that are too similar to their children will be discarded from the results. This helps to focus on the low level functions that are directly causing the spikes, since these are more likely the leaf functions where the CPU time is actually spent doing work.

5.2.2.1 Rare Button

The **Rare** button selects functions that rarely spike (which are sampled very infrequently).

Rare Button Options

The default is 10 or fewer spikes over the duration of the profile, but this can be changed with the **Max Spikes** option. The **Min Width** option controls how large a sample run must be (number of consecutive samples) in order to be considered a spike. You can adjust the **Min Width** to make the spike detection more or less sensitive.

5.2.2.2 Burst Button

The **Burst** button selects functions that burst irregularly (that tend not be called every frame).

Burst Button Options

The **Irregularity** option controls the threshold of frame-to-frame irregularity in order for a function to be detected (larger numbers represent more irregularity). The **Min Width** option controls how large a sample run must be (number of consecutive samples) in order to be considered a spike. You can adjust the **Min Width** to make the spike detection more or less sensitive.

5.2.2.3 Wild Button

The **Wild** button selects functions that wildly over their lifetime.

Wild Button Options

The first **Min Change** option controls the minimum number of total milliseconds a function must change over the entire profile in order to be detected (reduce this number to make it more sensitive). The second

Min Change option controls the minimum percentage a function must change over the entire profile in order to be detected (reduce this number to make it more sensitive).

5.2.2.4 Flutter Button

The **Flutter** button selects functions that are routinely called every frame that have a sudden increase or decrease from one frame to the next.

Flutter Button Options

The **Irregularity** option controls the threshold of frame-to-frame irregularity in order for a function to be detected (larger numbers represent more irregularity). The first **Min Change** option controls the minimum number of total milliseconds a function must change from frame-to-frame in order to be detected (reduce this number to make it more sensitive). The second **Min Change** option controls the minimum percentage a function must change frame-to-frame in order to be detected (reduce this number to make it more sensitive).

5.2.2.5 Swell Button

The **Swell** button selects functions that are routinely called every frame, but that gradually increase or decrease over time with a large total difference over the entire profile. For example, a function might take 5% on the first frame and then incrementally take more time each frame until by the last frame it takes 20%, but the change was gradual with no significant jumps in percentage from one frame to the next.

Swell Button Options

The **Irregularity** option controls the threshold of frame-to-frame irregularity in order for a function to be detected (larger numbers represent more irregularity). The first **Min Change** option controls the minimum number of total milliseconds a function must change over the entire profile in order to be detected (reduce this number to make it more sensitive). The second **Min Change** option controls the minimum percentage a function must change over the entire profile in order to be detected (reduce this number to make it more sensitive).

5.2.2.6 Bad Frame Button

The **Bad Frame** button selects the slowest frames and then selects functions that performed worse on those frames compared with all other frames.

Bad Frame Button Options

The **Max Frames** option controls how many slow frames it will select for the analysis. The **Use Selected Frames** option will use any user selected frames for the analysis, so you can focus on a particular frame or set of frames that you think are problematic (if no frames are selected, the **Bad Frame** analysis acts as if this option is disabled).

5.2.3 Quick Filters

The buttons in the **Quick Filters** section allow you to quickly understand a profile across several different metrics.

Everything done by the **Quick Filters** can also be done manually, however the **Quick Filters** allow a one-click solution to quickly analyzing a profile. When a quick filter button is pressed, it first resets the state of the GUI by performing the following actions:

- Deselects all functions.
- In the **Functions** tab: Selects **All Cores** and **All Threads**, clears the filter, and the functions are sorted by the **Self** heading.
- In the **Sample Graph** tab: Selects **Total** if **Callstack** profile, otherwise **Self** is selected.

Currently, all **Quick Filters** will perform the following steps:

1. In the **Functions** tab, the **Total** header is selected if it is a **Callstack** profile, otherwise the **Self** header is selected.
2. In the **Functions** tab, the filter text is set to a particular string with **Regex** enabled. This attempts to isolate all functions related to the particular button pressed. Later versions of the profiler will allow you to customize each quick filter text string, however you are free to alter the text filter in the **Functions** tab after you have clicked one of the **Quick Filters**.
3. In the **Sample Graph** tab, **Only Combined** is selected. This will cause all of the selected functions to be summed and drawn as a single load line for each core.
4. In the **Functions** tab, the **Select Top** button is triggered to select the top 45 functions.

Idle

The **Idle** filter shows idle time per frame on each core. All games will have some idle time, but excessive idle time indicates that a core isn't being well utilized.

Physics

The **Physics** filter shows physics time per frame on each core (including Havok and Bullet middleware).

GFX

The **GFX** filter shows graphics time per frame on each core. This filter also includes the **Particle** filter.

Particle

The **Particle** filter shows particle system time per frame on each core.

Anim

The **Anim** filter shows animation time per frame on each core.

AI

The **AI** filter shows artificial intelligence time per frame on each core (behavior and navigation).

Audio

The **Audio** filter shows audio time per frame on each core (including Audiokinetic, Ogg Vorbis, and FMOD middleware).

Online

The **Online** filter shows online time per frame on each core (including Quazal middleware).

Job Sys

The **Job Sys** filter shows job system time per frame on each core.

Math

The **Math** filter shows math time per frame on each core.

Mem

The **Mem** filter shows memory time per frame on each core.

Reset

The **Reset** quick filter button will perform the reset as described at the beginning of this section. Use this to restore the GUI to a default state after using the various **Quick Filters**. Most importantly, on the **Sample Graph**, it will change the **Only Combined** setting back to **Hide Combined**.

5.2.3.1 Customizing the Quick Filters

All of the Quick Filters are customizable, and it is possible to create new Quick Filters as well. In the **Directories** tab there is a **Button Dir** which can be set to a folder to attempt to load Quick Filters from it (see [Button Directory](#)).

Changing Behavior of an Existing Quick Filter

To change the behavior of the default Quick Filters, copy the original Quick Filter XML file from the Quick Filter folder (located in `bin\Quick Filters`) into your custom button directory. Now you can edit the behavior of the Quick Filter to suit your application. When the buttons are loaded your custom button will be loaded instead of the original, default version of the button.

Creating a New Quick Filter

To create a new Quick Filter, it is recommended that you start by copying one of the existing Quick Filter XML files into your custom button directory. You can then edit the fields of the button to get the behavior that you are looking for.

It is especially important to change the `<Title></Title>` field in the XML file. The Title field is used in order to determine uniqueness of the buttons. If the Title matches another Title, the new button will overwrite the original button.

Using System Classification

Quick Filter scripts have the ability to use the regular expressions located in the `Classification.xml` file. This allows for a Quick Filter to use the same search as the Classification feature (for more information Classifications, see [System Classifications](#)). To access these classifications in the script, use the standard Windows PowerShell variable markup, with the name of the variable the same as the XML tag that defines the expression.

When the Quick Filter script is loaded, it will perform a simple find-and-replace on the provided script against the classification values loaded. As such, it is strongly recommended that any scripts written use variable names of a different format than those used to define classifications. The default classifications all take the form of `$<System>Classification`, for example `$AudioClassification`.

5.3 Heartbeats Tab

The **Heartbeats** tab allows profile data to be framed against particular heartbeats per core. In addition, a subset of frames at a given rate can be quickly selected.

5.3.1 Group Selection

Group selection allows particular groups of frames to be selected according to their rate. The following table explains each button's action.

Table 5-1: Group Selection Buttons and Their Actions

Button	Action
60Hz	Select all frames on all cores running at 60 Hz or faster (16.6 ms or less)

Button	Action
30Hz	Select all frames on all cores running between 30 Hz and 60 Hz (16.6 – 33.3 ms)
20Hz	Select all frames on all cores running between 20 Hz and 30 Hz (33 – 50 ms)
15Hz	Select all frames on all cores running between 15 Hz and 20 Hz (50 – 66.6 ms)
12Hz	Select all frames on all cores running between 12 Hz and 15 Hz (66.6 – 83.3 ms)
Under 12Hz	Select all frames on all cores running lower than 12 Hz (greater than 83.3 ms)
Under 30Hz	Select all frames on all cores running lower than 30 Hz (greater than 33.3 ms)
Invert	Invert frame selection (selected frames are deselected and unselected frames are selected)
Deselect All	All frames are deselected

Selection of frames can be viewed in the **Sample Graph** tab. Note that all of the buttons can be toggled, except for **Invert** and **Deselect All**. If you toggle a particular button on, the corresponding frames that match the rate will be selected (highlighted in orange). Toggling a button off will deselect frames that match the rate. The toggling action allows you to select or deselect several ranges independently, for example allowing you to select frame rate ranges. You can manually override these group selections by directly selecting or deselecting individual frames within the **Sample Graph**.

When frames are selected or deselected, all displayed values (for example percentages in the **Functions** tab) are recalculated to only include the selected frames. If no frames are selected, the entire data set is reflected in the values.

5.3.2 Heartbeat Cores

For each core that was recorded during a profile, a numbered **Core** group will appear in the **Heartbeats** tab. Within that **Core** group, you can select heartbeat framing for that core along with some toggles to make that framing apply to all cores or as the master frame rate.

There are three types of selections that appear in the drop down box:

1. Fixed intervals (**Fixed 60Hz Intervals**, **Fixed 1ms Intervals**).
2. Heartbeats that were recorded directly by your game (for example **MainHeartbeat**).
3. Inferred heartbeats (periodic rhythms that were automatically detected in your game).

Fixed intervals are offered in order to draw fixed time markers within the **Sample Graph** for that core. By design, they can't be used as frame boundaries or to select frames, since the game's profile data would never be in-sync with these fixed intervals and the resulting information would be misleading.

Recorded heartbeats and inferred heartbeats show useful frame markings on the **Sample Graph** that can then be used for further analysis or selection. For example, you can select all **30Hz** frames from the **Group Selection** ribbon bar group. Recorded heartbeats are manually logged in the game code using the functions as described in the [Wii U CPU Profiler Game API](#). Inferred heartbeats are automatically detected, based on periodic function behavior. However, if the sampling rate is too low or there are no detectable patterns, then there may not be any inferred heartbeats. Also note that inferred heartbeats will have arbitrary frame boundaries based on the most periodic function detected, so it is unlikely that it will line up with the true beginning of your frame (we recommend that you deliberately record a heartbeat in your code if you want accurate frame boundaries). Lastly, inferred heartbeats are ordered from most represented in the profile to least represented, so pay more attention to inferred heartbeats that appear higher up within the drop down box.

If a recorded heartbeat or inferred heartbeat is selected, then that heartbeat can be used across all cores by checking the **Apply to All Cores** toggle button. Additionally, the currently selected heartbeat can be

used as the master frame rate (by checking the **Use as Frame Rate** toggle button), which will cause it to be drawn in the timeline at the bottom of the **Sample Graph**. By default, the recorded heartbeat (frame marker) is used as the master frame rate and drawn in the **Sample Graph** timeline.

5.4 Sampled Profile Tab

There are several settings that can affect a **Sampled Profile**. Please consider the following settings.

As the settings on this tab all relate to taking a profile, this tab is only visible when connected to a dev kit. When the connection is established, this tab will automatically appear in the **Ribbon Bar** and become the selected tab.

5.4.1 Profile Buttons

The **Start** and **Stop** profile buttons are available on the far left side. Once the **Start** button is pressed, it will change to a **Stop** button so that the profile can be manually stopped without moving the mouse cursor.

The number to the right of the **Start** button will time how long you have been taking a profile. It will begin counting when you hit the **Start** button.

The drop down box next to the **Start** button controls when the profile will stop. To manually stop the profile by pressing the **Stop** button, leave the drop down selected to **Manual Stop**. If you want to record for a set period of time, choose a time in the drop down box. You can still **Stop** a profile at any time before the specified amount of time has elapsed.

Profiling will always stop when the profile buffer is full.



Important: The **Start** button on the **Sampled Profile** tab will start a sampled profile (as opposed to an **Instrumented Profile**). If you want to start an **Instrumented Profile**, switch to the **Instrumented Profile** tab and use that **Start** button.

5.4.2 Sampling Strategy and Rate Buttons

This group of controls will determine the event that triggers a sample to be taken and the rate at which the event should trigger. In general, there are two types of events: **Sample by Time** and **Sample by Performance Counter**.

Sample by Time uses a timer (jittered for good sampling) to determine when to sample. **Sample by Performance Counter** samples every n counter events as specified by the **Sampling Rate** drop down box.

Sampling Strategy Drop Down Box

This control determines the event that triggers a sample to be taken. The following options are available for sampling strategies (where n is defined in the **Sampling Rate Drop Down Box**):

Sample by Time

The default sampling option. Samples are taken roughly n times per frame 60 Hz frame.

Sample by L1 Data Cache Miss

Samples are taken every n times that an L1 data cache miss occurs.

Sample by L2 Data Cache Miss

Samples are taken every n times that an L2 data cache miss occurs.

Sample by L1 Instruction Cache Misses

Samples are taken every n times that an L1 instruction cache miss occurs.

Sample by L2 Instruction Cache Misses

Samples are taken every n times that an L2 instruction cache miss occurs.

Sample by L1 Misses to MEM

Samples are taken every n times an L1 miss occurs and data is loaded from main memory rather than from L2 or via Intervention.

Sample by L1 Write to L2

Samples are taken every n times a write from L1 to L2 occurs. Note that storage related interrupts trigger when the store completes, so the instruction or function in which the sample is taken may be different than the one that initiated a store.

Sample by L2 Write to Memory

Samples are taken every n times a write from L2 to Main Memory occurs. Note that storage related interrupts trigger when the store completes, so the instruction or function in which the sample is taken may be different than the one that initiated a store.

Sample by Cycles Stalled L1 Miss

Samples are taken every n cycles that are counted when L1 misses occur.

Sample by Branch Misprediction

Samples are taken every n times a branch misprediction occurs.

Sampling Rate Drop Down Box

The profiler has several sample rates to choose from, based on the **Sampling Strategy** selected. If **Sample by Time** is selected, the rates reflect the number of times to sample during a 60 Hz frame. For example, a rate of **500x per frame** will result in 500 samples per 60 Hz frame. In the case that the game runs at 30 frames per second, a rate of 500x will effectively sample the game at double the stated rate, or 1000 times per frame. A game that runs at 20 frames per second will effectively be sampled at triple the stated rate, or 1500 times per frame. If one of the **Sample by Performance Counter** choices are selected, then the rate will be in terms of that performance counter.

When sampling by time, as the sampling rate is changed, the estimated profiling time is displayed. As sampling by performance counter does not use a timer no estimate of the profiling time can be given.

5.4.3 Data per Sample

This group of controls determines what is recorded with each sample.

The number of bytes per sample will be shown on the top left side of the **Data per Sample** group. This number is based on all of the **Data per Sample** settings selected.

5.4.3.1 Performance Counters Drop Down

The **Performance Counters** drop down box allows you to choose which performance counters will be recorded with each sample.

Please see [Performance Counter Groups](#) for more information on what performance counters are available and what kind of data they record.

5.4.3.2 Callstack Toggle Button

There are two choices for what kind of function data should be recorded with each sample:

Function: This is the option selected if the **Callstack** button is unchecked. When a sample event is triggered only the function that is currently executing is recorded. As very little data is recorded, this option will generally allow for much longer profile times.

Callstack: This is the option selected if the **Callstack** button is checked. When a sample event is triggered the entire callstack is recorded. This option records a variable amount of memory, depending on the depths of the callstacks recorded. Recording the callstack has the following implications:

- Good: The **Total** time a function has executed for (the time within the function and all functions it calls) will be shown in the **Functions** tab in addition to its **Self** time.
- Good: A call tree will be constructed from the individual profiles and displayed in the **Call Tree** tab. The call tree shows the multiple places where each function was called from and the amount of time spent in each function of the tree.
- Bad: The buffer is filled much more quickly, resulting in shorter possible profile times.
- Bad: There is much more overhead from the profiler, which hurts the frame rate and the accuracy of the profile (since cache is negatively affected by more profiler work).

The number of bytes per sample will be shown on the top left side of the **Data per Sample** group. This number is based on all of the **Data per Sample** settings selected.

5.4.3.3 OS Stats Button

The **OS Stats** button enables and disables the recording of OS statistics and Usage statistics in your profile data. These statistics are fairly large, so it is recommended that this feature be used sparingly.

Please see [OS Statistics](#) and [Usage Statistics](#) for more information on what kind of information the OS statistics can provide.

5.4.3.4 GPU Stats Button

The **GPU Stats** button enables and disables the recording of GPU statistics in your profile data.

Please see [GPU Statistics](#) for more information on what kind of information the GPU statistics can provide.

5.4.4 Cores

The buttons in this section turn on/off profiling for each core. Sometimes it is advantageous to only profile on a single core since it will allow significantly longer profile times and reduce profiler overhead on the other cores. There must be at least one core selected.

5.4.5 Recording Marked Code Blocks

To turn on recording of code blocks, select the **Marked** button. Note that code blocks must be manually instrumented in the game's source code using the [Wii U CPU Profiler Game API](#). This button simply turns on recording of them during a profile.



Caution: Recording code blocks can hurt the recorded profile results.

5.5 Instrumented Profile Tab

This tab is used to take an **Instrumented Profile**. There are three choices to make before taking an instrumented profile.

1. Choose a function to profile (optional if **Marked** is selected).
2. Choose a performance counter group to record.
3. Choose to have **Marked** code blocks recorded (optional if a function is chosen).

As the settings on this tab all relate to taking a profile, this tab is only visible when connected to a dev kit. When the connection is established, this tab will automatically appear in the **Ribbon Bar** and become selectable.

5.5.1 Allowing for Instrumentation

Certain conditions must be to allow for an instrumented profile to be taken.

In the latest SDKs, there are protection mechanisms to ensure that code is not modified at runtime. This will cause the profiler to crash when you attempt to instrument a function. In order to get around this, your application must be started in Debug Mode. In order to start you application in Debug Mode, use the following `caferun` command:

```
caferun -d multi -R RpxPath
```

The `-R` option keeps the MULTI debugger interface from starting until an `OSDebug` routine is called. Additional commands can be added as long as the two shown above are also included.

If you are using the SDK-provided makefiles, running in this manner can be accomplished by using the command:

```
make multirun BUILD_TYPE=NDEBUG USER_RUN=-R
```

5.5.2 Profile Buttons

The **Start** and **Stop** profile buttons are available on the far left side. Once the **Start** button is pressed, it will change to a **Stop** button so that the profile can be manually stopped without moving the mouse cursor.

The number to the right of the **Start** button will time how long you have been taking a profile. It will begin counting when you hit the **Start** button.

The drop down box next to the **Start** button controls when the profile will stop. To manually stop the profile by pressing the **Stop** button, leave the drop down selected to **Manual Stop**. If you want to record for a set period of time, choose a time in the drop down box. You can still **Stop** a profile at any time before the specified amount of time has elapsed.

Profiling will always stop when the profile buffer is full.



Important: The **Start** button on the **Instrumented Profile** tab will start an instrumented profile (as opposed to a **Sampled Profile**). If you want to start an **Sampled Profile**, switch to the **Sampled Profile** tab and use that **Start** button.

5.5.3 Choose a Function to Instrument

To choose a function to instrument, follow these steps:

1. Take a **Sampled Profile** so that the **Functions** tab is populated.
2. Click the **Add** button so that the button is highlighted.
3. Click on a function in the **Functions** tab, **Last** tab, **Call Tree** tab, or **Code Coverage** tab that you want to profile.

The function name should appear in the **Instrumented Profile** tab. When using the **Code Coverage** tab, it helps to use the filtering box at the bottom to find a particular function.

To choose a different function, click the **Add** button again and click on a different function in either the **Functions** tab or the **Code Coverage** tab. To remove your function from being selected, click on the red **X** button next to the function name in the **Instrumented Profile** tab.

5.5.3.1 Limitations on Selecting Functions To Instrument

The profiler does some basic error checking on the function that is selected. However, it is still possible to select a function that will break when instrumented. This stems from how the profiler instruments a function: it overwrites the first assembly instruction with a branch to our instrumentation code, does some work, executes the original instruction, and then returns the second instruction in the instrumented function. The following code structures are things to be aware of that the profiler does not currently check for automatically:

- Functions that branch back to the first instruction.

If the above are not occurring in the function but a problem arises, please contact your local Nintendo support group.

5.5.4 Choose a Performance Counter Group

You can record performance counters along with an instrumented function or along with code blocks. Simply select your desired group from the **Performance Counter Drop Down Box**.

Recording performance counter data will use up buffer space slightly faster and cause additional profiler overhead.

Please see [Performance Counter Groups](#) for more information on what performance counters are available and what kind of data they record.

5.5.5 Recording Marked Code Blocks

To turn on recording of code blocks, select the **Marked** button. Note that code blocks must be manually instrumented in the game's source code using the [Wii U CPU Profiler Game API](#). This button simply turns on recording of them during a profile.



Caution: Recording code blocks can hurt the recorded profile results.

5.6 Directories Tab

The **Directories** tab allows you to specify a variety of different important directories that the profiler will use to locate files.

Each of the directories can be edited in the following manner:

1. **Directly edit the text box:** You can type the directory in the text box. If the directory does not exist, the background of the text box will turn orange. Once a valid directory is entered, the background will return to normal. If an invalid directory is in the text box and focus leaves the text box, the directory will be replaced with the last known valid directory.
2. **Browse for the directory:** You can click the browse directory button to the right of the text box to find the desired directory.



Note: You can reset the directory back to default by directly editing the text box and clearing the contents.

5.6.1 SLC Directory

The **SLC Dir** option is used to specify where your SLC directory exists. This is important so that the profiler can find associated RPLs and such. When you specify this directory, it will be saved in the settings for the next time the profiler is loaded.

The SLC directory contains settings files and the system RPLs. The directory is one of the SDK subdirectories. Setting this directory allows for the profiler to find and load the system RPLs, thereby greatly reducing the number of UNRESOLVED symbols.

If the `CAFE_ROOT` environment variable is set, **SLC Dir** will default based on the `CAFE_ROOT` variable's location. If the directory is changed, that setting is saved and used from that point onward.

If the `CAFE_ROOT` environment variable is not set, **SLC Dir** will default to the `c:` drive.

5.6.2 Button Directory

The **Button Dir** is used to specify the folder from which custom **Quick Filter** buttons will be loaded. This allows the user to create custom buttons that can be easily preserved between Profiler versions. For information on how to create custom buttons see [Customizing the Quick Filters](#).

Button Dir will default to an empty box - meaning no custom buttons will be loaded.

5.6.3 Code Directory

The **Code Dir** is used to specify the folder which the **Assembly Tab** should use as an attempt to find your source code. This is useful if looking at a profile that was taken on someone else's machine which does not have a matching location for their code folder.

To find the source file, the profiler uses the GHS tools to turn specific addresses into filenames and line numbers. When a **Code Dir** is specified, the **Assembly Tab** searches for the code file it will start by looking for the base file name in the folder specified in **Code Dir**. If the file is not found it will continually add a directory from the base filename until it finds the file or runs out of directories to add.

If **Code Dir** is not specified, or if a search of the **Code Dir** came up empty, the **Assembly Tab** will default to trying to load in the file as returned by the GHS tool.

Code Dir will default to an empty box.

5.7 Quick Ribbon Bar

The Quick Ribbon Bar is located in the top-right corner of the profiler. It provides a quick way to perform the following actions, even when the Ribbon Bar is minimized.

- **Sync** to a dev kit.
- **Unsync** from a dev kit.
- **Start a Sampled Profile**.
- **Stop** taking a profile.
- **Save** the current profile to disc.

6 Functions Tab

The **Functions** tab displays all of the functions that were sampled during a profile.

6.1 Sorting Functions

Functions can be sorted by **Total**, **Self**, **Core**, **Thread**, **Module**, **Opcodes**, or **Name** by clicking on each respective heading. Please note that the **Total** column will only be shown if a **Callstack** profile was taken (see [Callstack Toggle Button](#)). The headings **Thread**, **Module**, and **Opcodes** will only appear when toggled on by clicking the **Thread**, **Module**, and **Ops** buttons, respectively.

By default, the functions are sorted by **Self**, which is the amount of time spent in each function (not including any sub-functions). Click on the **Total** heading to sort by **Total**, which is the amount of time spent in each function and every function it calls.

6.2 Functions Tab Toolbar

The following controls are in the **Functions** tab toolbar.

6.2.1 Core Filtering and Column Heading

If a profile contains multiple cores, you can click on the **Cores** drop-down in the toolbar to show a list allowing for selection between **All Cores** and each individual core. Only cores that were profiled will be available in the **Cores** drop-down. The default selection is **All Cores** and each new profile taken or loaded will reset the **Cores** drop-down selection to **All Cores**.

When a particular core is selected, such as **Core 1**, the **Core** column will be hidden since all functions shown belong to the same core. This can be helpful when focusing on a particular core and can help optimize horizontal space in the tab.

6.2.2 Deselect All Button

To quickly deselect all selected frames or selected periods of time, click on the **Deselect All** button.

6.2.3 Select Top Button

To quickly select the top 45 functions, click on the button marked **Select Top** in the **Functions** toolbar.

6.2.4 Diff Total and Diff Self Column Headings

By default, the **Diff Total** and **Diff Self** columns are hidden (to conserve horizontal space), but can be enabled by clicking on the **Diff** button in the **Functions** toolbar. Click the button again to hide the columns.

The **Diff Total** and **Diff Self** columns compare the current profile with the last loaded profile, showing a percent difference for each function. For example, if a particular function's **Self** percentage is 10% in the last loaded profile and 12% in the current profile, then a value of +2% will be shown in the **Diff Self** column. Likewise, if a particular function's **Self** percentage is 10% in the last loaded profile and 8% in the current profile, then a value of -2% will be shown in the **Diff Self** column.

To see how your game's performance has changed over time, you can first load a profile taken a while ago (i.e. yesterday, a week ago, a month ago), and then take a current profile. Within the **Diff Total** and **Diff Self** columns it will show how much each function has increased or decreased as a percentage. However, be aware that if a particular function was optimized to take less time (and its percentage decreased), percentage-wise it will cause all other functions to increase (even if the absolute amount of time they

took was the same between both profiles). To sort all functions by the absolute value of the percentage difference, click on either the **Diff Total** or **Diff Self** column headers. The sorting is based on absolute value in order to prioritize functions that had the largest percentage change, whether that change was positive or negative.



Note: The **Diff Total** column is only available on callstack profiles.

6.2.5 Threading Filtering and Thread Column Heading

By default, the **Thread** column is hidden (to conserve horizontal space). However, if you want to see which thread a function belongs to, click the **Thread** button in the toolbar. Click the button again to hide the column.

If you want to only view functions belonging to a particular thread, click the **Thread** drop-down button and select a particular thread. When an individual thread is selected, the **Thread** column will be hidden since all shown functions belong to the same thread. The default selection is **All Threads** and each new profile taken or loaded will reset the **Thread** drop-down selection to **All Threads**.

By default, threads are named with a 32-bit identifier (such as 0x1052D298). If the profiler is aware of a name for the thread, that name will be displayed in the drop-down. However, the **Thread** column will always only use the 32-bit identifier.

6.2.6 Module Column Heading

By default, the **Module** column is hidden (to conserve horizontal space), but can be enabled by clicking on the **Module** button in the **Functions** toolbar. Click the button again to hide the column.

The **Module** column contains the name of the module that each function came from. To sort all functions by the module they came from (alphabetical sort), click on the **Module** column header.



Note: The **Module** column is only useful if there were relocatable modules used in the application.

6.2.7 Opcodes Column Heading

By default, the **Opcodes** column is hidden (to conserve horizontal space), but can be enabled by clicking on the **Ops** button in the **Functions** toolbar. Click the button again to hide the column.

This column shows how many assembly opcodes are in each function (the size of the function). To sort all functions by the number of assembly opcodes (most to least), click on the **Opcodes** column header.

6.2.8 Args Button

By default, function names don't include their argument list when seen in the **Functions** tab and **Call Tree** tab (in order to make the list more readable and speed up GUI drawing time). If you wish to show the function arguments, click the **Args** button.

6.2.9 Copy Button

The **Copy** button will copy to the clipboard whatever is currently shown using the current sorting and filtering. You can then paste the content into an email, document, or Excel spreadsheet. Because the copied content uses tab delimiters, it will retain the column organization when pasted into an Excel spreadsheet.

The **Copy** button takes into account what is currently selected. If any items are selected at the time the **Copy** button is pressed, then only the selected items will be copied. If nothing is currently selected, everything will be copied.

6.2.10 System Classification Filter

If you want to only view functions belonging to a particular system (like Graphics, Physics, Animation, etc.), click the **System Classification** drop-down button and select a particular system. The default selection is **All Systems** and each new profile taken or loaded will reset the **System Classification** drop-down selection to **All Systems**.

6.2.11 Selected Only Button

The **Selected Only** button forces the list display to only show items which are currently selected. If this mode is enabled and an item is deselected, it will be removed from the displayed list.

6.3 Filtering Functions

Functions can be filtered using the **Filter** text box at the bottom of the **Functions** tab. This allows you to find a particular function or if you want to look at a class of functions or a certain namespace (such as GX2).

Type any string into the text box and a case-insensitive search is performed for the items that contain the string. Only items matching the search text will be displayed. To negate the results (show everything that does not match the filter), click on the exclamation point button to the right of the text box. If there are no matches for the search, the text box will be highlighted orange.

In order to improve the filtering results, only the function name up to and including the first left parenthesis is considered. This eliminates results where a match occurs in the function arguments, since showing matches to function arguments is typically not desirable and results in unwanted matches.

The following are some helpful filters you can use:

Table 6-1: Helpful Simple Filters

Filter Text	! Btn	Description
(	N	Show all C++ functions.
(	Y	Show all C functions.
;	N	Show all functions that share code with other functions.

Regular expressions are also supported (using .NET RegEx, which is Perl 5 compatible). To have the filter interpreted as a regular expression, click on the **RegEx** button to the right of the text box. Unlike basic filters, regular expressions are case-sensitive.

The following are some helpful regular expressions:

Table 6-2: Helpful Regular Expression Filters

Filter Text	! Btn	Description
^GX2	N	Show all functions starting with "GX2".
^GX2	Y	Show all functions not starting with "GX2".
^GX2 ^__	Y	Remove multiple functions (not starting with "GX2" or "__").
(?i) ^cos ^sin ^tan	N	Show all trigonometric functions (case-insensitive for all).
(?i) a(?-i)nim	N	Show functions with either "anim" or "Anim".
(?i)math.*vec	N	Show functions with "math" and "vec" in that order (case-insensitive).

Filter Text	! Btn	Description
<code>(?<!trans)pose</code>	N	Show functions with "pose" which are not preceded by "trans".
<code>skin(?!g)</code>	N	Show functions with "skin" which are not followed by "g".



Note: All filters in the **Functions**, **Call Tree**, and **Code Coverage** tabs act similarly on functions and share saved filters.

6.3.1 Quick Regular Expression Primer

Regular expressions are a concise method for matching strings of text. The following are some rules that you might find helpful when specifying a regular expression:

- There should be no spaces in a regular expression (unless you're searching for a space).
- A vertical bar indicates Boolean "or". For example, `cos|sin` can match "cos" or "sin".
- `(?i)` indicates case-insensitive. Using it once causes everything afterward to be affected. For example, `(?i)cos|sin` matches "cos" or "sin" with both being case-insensitive.
- A caret indicates the beginning of a function only. For example, `^GX2` matches functions starting with "GX2".
- A period can match any character. An asterisk specifies zero or more of the preceding element. Together they can specify any character zero or more times. For example, `math.*vec` matches "math" and "vec" with zero or more characters between them.

6.3.2 Saved Filters

Saved filters allow for using the same filters between profiling sessions, with any saved filters being loaded when the profiler is opened. By default, the **Filter Drop Down** box is loaded with example filters that you can customize. If you want to save your own filter in the drop down box, press the **Save** ("+") button after typing the filter (the **Negate** and **RegEx** button states will also be saved with the filter text). To delete a filter from the drop down box, select it first, then press the **Delete** ("-") button.



Note: The saved filters in the **Functions** tab will also be available in the **Call Tree** and **Code Coverage** tabs. However, please note that the **Filter Drop Down** box in the **Threads** tab and **Counters** tab are not shared since they are used to search for thread and counter names, not function names.

6.4 Selecting Functions

Click on a function to select it and click again to unselect it. Up to 500 functions can be selected (each will be highlighted in a different color). Selected functions will be graphed in the **Sample Graph** and highlighted in the **Call Tree** (the **Call Tree** is only visible if the profile was taken with **Callstack** profile data - see [Callstack Toggle Button](#)).

To quickly select the top 50 functions, click on the button marked **Select Top** in the **Functions** toolbar.

To quickly deselect all functions, click on the **Deselect All** button in the **Functions** toolbar.

6.4.1 Alternative Selection Methods

Multiple items can be selected or deselected with Shift+Click (select/deselect an item without holding the Shift key, then hold the Shift key and click on a second item). The range between the two clicks will be either selected or deselected.

To quickly select only a single item and deselect everything else, Alt+Click can be used. This deselects all and then selects the clicked upon item.

6.5 Right-Click Context Menu

A right-click in the tab will bring up a context menu containing a variety of options.

6.5.1 Copy Function Name to Clipboard

This will copy the full function name with arguments to the clipboard.

6.5.2 Deselect All Other Functions

The Deselect All Other Functions option will force all functions other than the one click on to become deselected. This will not result in an unselected function becoming selected.

6.5.3 Show Assembly

The **Show Assembly** option will move focus to the [Assembly Tab](#), showing the assembly of the clicked function.

6.5.4 Show Details

The **Show Details** option will open a new [Function Details Window](#), showing the details of the clicked function.

6.6 Units

To change the units displayed before each item, use the **Unit** buttons in the **Home** tab of the ribbon bar (see [Units](#)).

6.7 Resizing the Window

All tabs in the left tab group, including the **Functions**, **Threads**, **Instrumented**, and **Counters** tabs, can be stretched horizontally by dragging the right edge with the mouse cursor. All of these tabs share the same display area, so resizing the width of one tab will result in all tabs in the tab group being resized.

6.8 Reading Obscured Function Names

If a function name is very long, it might be obscured by the right edge of **Functions** tab. You can either resize the **Functions** tab or you can hover over the function and a tooltip containing the full name will be displayed. Note that the tooltip will contain the full function signature including function arguments, regardless of the state of the **Args** button.

6.9 Multiple Function Names on a Single Line

Some lines will be a concatenation of several function names, each separated by a semi-colon (only visible when the **Args** button is selected or by hovering over the function and viewing the tooltip). For example, you might see the following function name listed on a single line in the **Functions** tab:

```
function_name_1; function_name_2
```

This happens when different functions get compiled into the same exact code. If the profiler samples execution in one of these shared functions, it can't determine which function name was called, so it concatenates all function names together on the same line.

6.10 RPL Internal Functions

When loading function names from an RPL, only the exported function names are visible.

Other functions are displayed as being an internal function. This display name takes the form

`filename.rpl_internal_xx` where '`filename.rpl`' is the name of the RPL from which functions were loaded and '`xx`' is a unique number for each individual function.

Each internal function is assigned its unique number based on its location in the RPL. As such, different builds of an RPL may have different numberings - so profile data cannot necessarily be compared when different RPLs are in use (for example, DEBUG vs. RELEASE or different SDK versions).

6.11 Unresolved Functions

When a profile contains a function it cannot resolve to a name, it will label the function as an internal unresolved function. These generally take the form of `filename.ext_internal (unresolved)` where

'`filename.ext`' is the name of the file from which function symbols were loaded.

1. The most likely cause for an unresolved function is that its name could not be found in the RPL. With changes to SDK 2.08 to decrease the number of functions in the symbol table, the profiler is no longer able to determine individual internal functions for NDEBUG builds of the system RPLs. All previous 'internal' functions are now marked as a single unresolved function for these NDEBUG RPLs.
2. The profiler could not locate the RPL or ELF that contained the function address. If the profiler could determine which RPL or ELF the function should have come from, the function will be listed as coming from '*RplName* (unresolved)'.

SDK RPLs are located by setting the **SLC Dir** value in the **Directories** tab.

7 Last Tab

The **Last** tab displays all of the functions from the last sampled profile. Not only is this useful for comparison purposes, but it is extremely convenient for selecting new functions to be instrumented, since once you take an **Instrumented Profile**, the normal **Functions** tab is hidden since it no longer contains data. The functions in the **Last** tab cannot be selected, but most of the operations that work in the **Functions** tab also work in this tab, such as all of the sorting and filtering functionality. Note that the units for function data on the **Last** tab are restricted to percentage and samples.

8 Threads Tab

The **Threads** tab displays all of the threads that were sampled during a profile.

8.1 Sorting Threads

Threads can be sorted by **Self**, **Core**, or **Name** by clicking on each respective heading. By default, the threads are sorted by **Core**, and then sorted by **Self** within each core.

8.2 Threads Tab Toolbar

The following controls are in the **Threads** tab toolbar.

8.2.1 Core Filtering and Column Heading

If a profile contains multiple cores, you can click on the **Cores** drop-down in the toolbar to show a list allowing for selection between **All Cores** and each individual core. Only cores that were profiled will be available in the **Cores** drop-down. The default selection is **All Cores** and each new profile taken or loaded will reset the **Cores** drop-down selection to **All Cores**.

When a particular core is selected, such as **Core 1**, the **Core** column will be hidden since all functions shown belong to the same core. This can be helpful when focusing on a particular core and can help optimize horizontal space in the tab.

8.2.2 Deselect All Button

To quickly deselect all selected frames or selected periods of time, click on the **Deselect All** button.

8.2.3 Top Button

The **Top** button selects the first 45 items in the list, depending on the currently selected sorting mechanism.

8.2.4 Copy Button

The **Copy** button will copy to the clipboard whatever is currently shown using the current sorting and filtering. You can then paste the content into an email, document, or Excel spreadsheet. Because the copied content uses tab delimiters, it will retain the column organization when pasted into an Excel spreadsheet.

The **Copy** button takes into account what is currently selected. If any items are selected at the time the **Copy** button is pressed, then only the selected items will be copied. If nothing is currently selected, everything will be copied.

8.3 Selecting Threads

Click on a thread to select it and click again to unselect it. Up to 45 threads can be selected (each will be highlighted in a different color). Selected threads will be graphed in the **Sample Graph**.

To quickly select the top 45 threads, click on the button marked **Top** in the **Threads** toolbar.

To quickly deselect all threads, click on the **Deselect All** button in the **Threads** toolbar.

8.3.1 Alternative Selection Methods

Multiple items can be selected or deselected with Shift+Click (select/deselect an item without holding the Shift key, then hold the Shift key and click on a second item). The range between the two clicks will be either selected or deselected.

To quickly select only a single item and deselect everything else, Alt+Click can be used. This deselects all and then selects the clicked upon item.

8.4 Filtering Threads

Threads can be filtered using the **Filter** text box at the bottom of the **Threads** tab. This allows you to find a particular thread when a large number of threads are shown. Typically there aren't that many threads, so this is usually unnecessary. It is usually better to filter by core using the **Core** drop down box in the **Threads** toolbar.

Type any string into the text box and a case-insensitive search is performed for the items that contain the string. Only items matching the search text will be displayed. To negate the results (show everything that does not match the filter), click on the exclamation point button to the right of the text box. If there are no matches for the search, the text box will be highlighted orange.

Regular expressions are also supported (using .NET RegEx, which is Perl 5 compatible). To have the filter interpreted as a regular expression, click on the **RegEx** button to the right of the text box. Unlike basic filters, regular expressions are case-sensitive.

Please see the [Quick Regular Expression Primer](#) for more information on Regular Expressions.

8.4.1 Saved Filters

Saved filters allow for using the same filters between profiling sessions, with any saved filters being loaded when the profiler is opened. By default, the **Filter Drop Down** box is loaded with example filters that you can customize. If you want to save your own filter in the drop down box, press the **Save** ("+") button after typing the filter (the **Negate** and **RegEx** button states will also be saved with the filter text). To delete a filter from the drop down box, select it first, then press the **Delete** ("-") button.



Note: The saved filters in the **Threads** tab are saved independently of all other filters.

8.5 Resizing the Window

All tabs in the left tab group, including the **Functions**, **Threads**, **Instrumented**, and **Counters** tabs, can be stretched horizontally by dragging the right edge with the mouse cursor. All of these tabs share the same display area, so resizing the width of one tab will result in all tabs in the tab group being resized.

9 Instrumented Tab

The **Instrumented** tab displays all instrumented functions and all instrumented code blocks that were recorded during a profile. The **Instrumented** tab will only appear if either of these types of data were recorded.

9.1 Instrumented Tab Toolbar

The following controls are in the **Instrumented** tab toolbar.

9.1.1 Deselect All Button

To quickly deselect all selected frames or selected periods of time, click on the **Deselect All** button.

9.1.2 Expand All Button

The **Expand All** button expands all possible items in the view.

9.1.3 Collapse All Button

The **Collapse All** button collapses all possible items in the view.

9.1.4 Select Similar Button

When the **Select Similar** button is enabled selecting or deselecting an item causes similar items to be selected or deselected.

On the **Instrumented** tab, the **Select Similar** button will allow for selecting similar values across instrumentation and code blocks. Only expanded blocks are evaluated for **Select Similar**. For example, a single click on one block's call count will result in the call count for all expanded blocks becoming selected (or deselected).

9.1.5 Copy Button

The **Copy** button will copy to the clipboard whatever is currently shown using the current sorting and filtering. You can then paste the content into an email, document, or Excel spreadsheet. Because the copied content uses tab delimiters, it will retain the column organization when pasted into an Excel spreadsheet.

The **Copy** button takes into account what is currently selected. If any items are selected at the time the **Copy** button is pressed, then only the selected items will be copied. If nothing is currently selected, everything will be copied.

9.2 Instrumented Items

Instrumented items can be expanded to view all of the data that was recorded during the profile. The following measurements should appear for each instrumented item.

- **Function Calls or Blocks:** The number of times the function or code block was entered and exited during the profile. If these are not the same number, the lesser of the two is used (for example, sampling stops while in a recursive function).
- **Max Recursive Depth:** Displays the max recursive depth that was detected during the profile. This measurement only appears if the value is greater than zero.

- **Time per Call or Time per Block:** This is the average time per call or block that was measured (time = total time / number of calls).
- **Performance Counter Data:** These measurements are based on the performance counters that were selected at the time of the profile. These values are over the duration of the profile. These values might be more useful if they are divided by the number of calls (for values that aren't percentages).

9.3 Selecting Instrumented Items to Graph

Selecting data fields within each instrumented item will cause them to be graphed within the **Sample Graph**. Up to 45 data fields can be selected at a time.

To select similar data fields within all expanded items, enable the **Select Similar** button in the **Instrumented** tab toolbar and click on a single data field.

9.3.1 Alternative Selection Methods

Multiple items can be selected or deselected with Shift+Click (select/deselect an item without holding the Shift key, then hold the Shift key and click on a second item). The range between the two clicks will be either selected or deselected.

To quickly select only a single item and deselect everything else, Alt+Click can be used. This deselects all and then selects the clicked upon item.

9.4 Resizing the Window

All tabs in the left tab group, including the **Functions**, **Threads**, **Instrumented**, and **Counters** tabs, can be stretched horizontally by dragging the right edge with the mouse cursor. All of these tabs share the same display area, so resizing the width of one tab will result in all tabs in the tab group being resized.

10 Counters Tab

The **Counters** tab displays all of the counters that were sampled during a profile.

10.1 Types of Counters

The profiler currently supports the following general counters. An individual platform may contain more counters.

10.1.1 Performance Counters

For more information on Performance Counters, please see [Performance Counter Groups](#).

10.1.2 OS Statistics

This option records some information about the various interrupts that occur in the OS while your game is active.

The currently provided OS Statistics are:

- All Interrupts
- GPU Interrupts
- Audio Interrupts
- Other Interrupts

10.1.3 Usage Statistics

This option records some basic information related to thread usage in your title.

The currently provided Usage Statistics are:

- Core Idle Time Percentage
- Thread Switch Counts

10.1.4 GPU Statistics

This option records various busy state information about the GPU at the time that sampling occurred.

The currently provided GPU Statistics are:

- 3D Engine Busy
- CB Busy
- CP Busy
- CPU Building Command List
- DB Busy
- DMAE Active
- DMAE Busy
- DMAE Waiting on Semaphore
- GPU Processing Command List
- GPU Waiting on Semaphore
- IB1 Buffer Size
- IB2 Buffer Size
- ME Parsing Packets
- PA Busy

- Pending Command Buffers
- PFP Parsing Packets
- SPI Busy
- SQ Busy
- Stream Out Busy
- TEX Busy
- VGT Busy

Following are more detailed descriptions of the counters that are shown. Most of the data below is planned to be added to the GX documentation in the future.

10.1.4.1 3D Engine Busy

This value will show whether any part of the GPU pipeline was busy when the sample was taken. This helps provide information over the frame on where you may be able to provide the GPU more work to perform.

10.1.4.2 CB Busy

The Render Backend (RB) is writing color data to memory.

10.1.4.3 CP Busy

The Command Processor (CP) is busy. Any of the follow can cause the CP to show as busy:

- Parsing commands
- Handling state management of registers
- Handling surface synchronization and memory semaphores
- Waiting for feedback (such as timestamps)

Because 'Waiting' time will show the CP being busy, this value can be initially misleading. We are working on finding a way to provide more details about when the CP is specifically waiting for feedback.

10.1.4.4 CPU Building Command List

This line will show when the CPU was building the command list that will later be processed by the GPU. In order for this to contain valid data you must call both `PROFILERGraphicsFrameBegin()` and `PROFILERGraphicsFrameEnd()` as directed in the API man pages.

Internally the profiler uses 4 buffers to track CPU command list building and GPU command list processing so that it can deal with buffered frames. The height of the line shows which of the 4 buffers was used. This can then be checked against the **GPU Processing Command List** line to see when the corresponding command list was processed by the GPU.

10.1.4.5 DB Busy

The Render Backend (RB) is writing depth data to memory.

10.1.4.6 DMAE Active

This line will show when the DMA Engine is either reading or writing memory.

10.1.4.7 DMAE Busy

This line will show when the DMA Engine is considered busy. This can be because a transfer is actively happening or if it is waiting for a semaphore.

10.1.4.8 DMAE Waiting on Semaphore

This line will show when the DMA Engine is waiting for the GPU to release a semaphore.

10.1.4.9 GPU Processing Command List

This line will show when the GPU was processing the command list that was created previously on the CPU. In order for this to contain valid data you must call both `PROFILERGraphicsFrameBegin()` and `PROFILERGraphicsFrameEnd()` as directed in the API man pages.

Internally the profiler uses 4 buffers to track CPU command list building and GPU command list processing so that it can deal with buffered frames. The height of the line shows which of the 4 buffers was processed by the GPU. This can then be checked against the **CPU Building Command List** line to see when the command list was built.

10.1.4.10 GPU Waiting on Semaphore

This line will show when the GPU is waiting for the DMA Engine to release a semaphore.

10.1.4.11 IB1 Buffer Size

This value shows how much data is remaining to be parsed for Indirect Buffer 1. If there is data remaining to be parsed, the Indirect Buffer is considered Busy. If there are long periods of non-zero value, this could be an indicator of a section where a large amount of work is being performed. Use Spark to gather more information on these sections.

10.1.4.12 IB2 Buffer Size

This value shows how much data is remaining to be parsed for Indirect Buffer 2. If there is data remaining to be parsed, the Indirect Buffer is considered Busy. If there are long periods of non-zero value, this could be an indicator of a section where a large amount of work is being performed. Use Spark to gather more information on these sections.

10.1.4.13 ME Parsing Packets

This line shows when the MicroEngine (ME) is actively parsing command packets.

10.1.4.14 PA Busy

The Primitive Assembly is busy. Any of the following can cause the PA to show as busy:

- Clip plane testing (trivial reject/accept)
- Clip space to Screen Coordinate transforms
- Clipping to frustum and user-defined planes
- Back-face culling
- Calculation of barycentric coordinates

10.1.4.15 Pending Command Buffers

This line shows when there are pending Indirect Buffer 1, Indirect Buffer 2, Command Buffer, or State Data requests that have not yet been initiated by one of the fetchers.

10.1.4.16 SPI Busy

At least one of the Shader Parameter Interpolator pipes is busy performing work. Any of the following can cause a busy state:

- Attribute interpolation for primitive pixels
- Loading input GPRs
- Selection of SIMD used for shader execution

10.1.4.17 SQ Busy

The sequencer is busy performing a task. Any of the following can cause a busy state:

- Shader instruction fetching and parsing

- Shader constant fetching

10.1.4.18 Stream Out Busy

The Stream Out feature is currently in use.

10.1.4.19 TEX Busy

At least one of the Texture pipes is busy performing work. Any of the following can cause a busy state:

- Computing texture coordinates and mip levels
- Texture filtering (ie, bilinear, trilinear, and anisotropic)

10.1.4.20 VGT Busy

The Vertex Grouper Tessellator is busy. Any of the following can cause a busy state:

- Vertex index data memory fetch
- Grouping index data into primitives
- Determination if a previous vertex can be re-used from the previous primitive
- Fixed function tessellation

10.1.5 User Logs

Custom recorded data using the profiler API.

Whenever the `PROFILERLog*Event` APIs are used, the data that is recorded will show up in the **Counters Tab** with the **Group** and **Name** specified in the `PROFILERLog*Event` call.

10.2 Sorting Counters

Counters can be sorted by **Core**, **Group**, or **Name** by clicking on each respective heading. By default, the counters are sorted by **Core**.

10.3 Counters Tab Toolbar

The following controls are in the **Counters** tab toolbar.

10.3.1 Core Filtering and Column Heading

If a profile contains multiple cores, you can click on the **Cores** drop-down in the toolbar to show a list allowing for selection between **All Cores** and each individual core. Only cores that were profiled will be available in the **Cores** drop-down. The default selection is **All Cores** and each new profile taken or loaded will reset the **Cores** drop-down selection to **All Cores**.

When a particular core is selected, such as **Core 1**, the **Core** column will be hidden since all functions shown belong to the same core. This can be helpful when focusing on a particular core and can help optimize horizontal space in the tab.

10.3.2 Group Filtering and Column Heading

The **Groups** drop-down in the **Counters** tab allows for selection between **All Groups** and any of the groups that were present in the profile (such as **Performance Counters**). Only groups that are in the profile will be listed in the **Groups** drop-down. The default selection is **All Groups** and each new profile taken or loaded will reset the **Groups** drop-down selection to **All Groups**.

When a particular group is selected, for example **Performance Counters**, the **Group** column will be hidden since all groups shown are of the same type.

Specifically for the group **Performance Counters**, the drop-down will contain additional info about which performance counter group was recorded for the profile (see [Performance Counter Groups](#)). For example, if a performance counter group named **Example Group** was profiled, the **Groups** drop-down will contain an entry for **Performance Counters (Example Group)**.

10.3.3 Deselect All Button

To quickly deselect all selected frames or selected periods of time, click on the **Deselect All** button.

10.3.4 Select Similar Button

When the **Select Similar** button is enabled selecting or deselecting an item causes similar items to be selected or deselected.

On the **Counters** tab, the **Select Similar** button will allow for selecting similar counters across cores. For example, if a performance counter group was profiled, a single click on the first counter on core 0 will select (or deselect) the first counter on all profiled cores.

10.3.5 Copy Button

The **Copy** button will copy to the clipboard whatever is currently shown using the current sorting and filtering. You can then paste the content into an email, document, or Excel spreadsheet. Because the copied content uses tab delimiters, it will retain the column organization when pasted into an Excel spreadsheet.

The **Copy** button takes into account what is currently selected. If any items are selected at the time the **Copy** button is pressed, then only the selected items will be copied. If nothing is currently selected, everything will be copied.

10.4 Filtering Counters

Type any string into the text box and a case-insensitive search is performed for the items that contain the string. Only items matching the search text will be displayed. To negate the results (show everything that does not match the filter), click on the exclamation point button to the right of the text box. If there are no matches for the search, the text box will be highlighted orange.

Regular expressions are also supported (using .NET RegEx, which is Perl 5 compatible). To have the filter interpreted as a regular expression, click on the **RegEx** button to the right of the text box. Unlike basic filters, regular expressions are case-sensitive.

Please see the [Quick Regular Expression Primer](#) for more information on regular expressions.

10.4.1 Saved Filters

Saved filters allow for using the same filters between profiling sessions, with any saved filters being loaded when the profiler is opened. By default, the **Filter Drop Down** box is loaded with example filters that you can customize. If you want to save your own filter in the drop down box, press the **Save** ("+") button after typing the filter (the **Negate** and **RegEx** button states will also be saved with the filter text). To delete a filter from the drop down box, select it first, then press the **Delete** ("-") button.



Note: The saved filters in the **Counters** tab are saved independently of all other filters.

10.5 Selecting Counters

Click on a counter to select it and click again to unselect it. Up to 45 counters can be selected (each will be highlighted in a different color). Selected counters will be graphed in the **Sample Graph** tab.

To quickly deselect all functions, click on the **Deselect All** button.

10.5.1 Alternative Selection Methods

Multiple items can be selected or deselected with Shift+Click (select/deselect an item without holding the Shift key, then hold the Shift key and click on a second item). The range between the two clicks will be either selected or deselected.

To quickly select only a single item and deselect everything else, Alt+Click can be used. This deselects all and then selects the clicked upon item.

10.6 Resizing the Window

All tabs in the left tab group, including the **Functions**, **Threads**, **Instrumented**, and **Counters** tabs, can be stretched horizontally by dragging the right edge with the mouse cursor. All of these tabs share the same display area, so resizing the width of one tab will result in all tabs in the tab group being resized.

10.7 Units

To change the units displayed before each item, use the **Unit** buttons in the **Home** tab of the ribbon bar (see [Units](#)).

11 Info Tab

The **Info** tab displays general information about the profile.

11.1 Info Tab Toolbar

The following controls are in the **Info** tab toolbar.

11.1.1 Expand All Button

The **Expand All** button expands all possible items in the view.

11.1.2 Collapse All Button

The **Collapse All** button collapses all possible items in the view.

11.1.3 Copy Button

The **Copy** button will copy to the clipboard whatever is currently shown using the current sorting and filtering. You can then paste the content into an email, document, or Excel spreadsheet. Because the copied content uses tab delimiters, it will retain the column organization when pasted into an Excel spreadsheet.

The **Copy** button takes into account what is currently selected. If any items are selected at the time the **Copy** button is pressed, then only the selected items will be copied. If nothing is currently selected, everything will be copied.

11.2 Environment

The **Environment** section has information on the profiler's operating environment at the time of the profile. This includes the following information:

- **User:** The name of the user logged into the computer.
- **Computer:** The name of the computer running the profiler.
- **Time:** The time at which the profile was taken.
- **Game:** The path to the executable file on the hard drive.
- **Captured By:** The name and type of kit used to capture the profile.

11.3 Sampling Settings

The **Sample Settings** section shows the settings that were used to generate the profile. These are listed on a per-core basis.

- The rate at which sampling was requested, either by time or by performance counter event trigger.
- Whether callstack profiling was enabled.
- Whether performance counters were recorded with the profile, and if so, which performance counter group was selected.
- Whether frame markers were seen on the core.

Example Output

```
Core 0: 10x per frame, Callstack Sampling, Has frame markers
```

11.4 Profile Statistics

The **Profile Statistics** section shows general information about the profile. This includes the following:

- **Buffer Used:** The size of the raw data buffer.
- Listed per Core
 - **Samples:** The number of samples seen on the core.
 - **Duration:** The duration of the profile on the core. It is possible for there to be a slight variance between core durations.
 - **Functions seen:** The number of functions seen executing on the core.
 - **Threads seen:** The number of threads seen executing on the core.
 - **Max callstack depth:** The deepest the callstack got at any time while profiling. (Only visible if a callstack profile was taken.)
 - **Call tree node:** The total number of nodes in the **Call Tree** tab. (Only visible if a callstack profile was taken.)

11.5 Module Statistics

The **Module Statistics** section shows the various modules that were seen during the profile. This is split into **Recorded Modules** and **Resolved Modules**.

Recorded Modules are those relocatable modules that the profiler saw loaded and/or unloaded. For each recorded module there should be an entry in the resolved modules list.

Resolved Modules are the actual modules that were loaded into the profiler in order to resolve symbol address and names. For static modules there will be an entry in this list, but not in the recorded modules list.

Each module listing may show the following:

- The name of the module.
- The path which the module was loaded from (only for **Resolved Modules**).
- **Base:** The address at which the module was loaded into memory.
- **Size:** The reported size of the module in memory (only for **Recorded Modules**).
- **Functions:** The number of functions loaded from the module (only for **Resolved Modules**).
- **Lifetime:** The times during which the module was loaded into memory in the format "[time loaded, time unloaded]".

11.6 Inferred Heartbeats

The **Inferred Heartbeats** section shows any inferred heartbeats that were found while analyzing the profile data. The following information is provided per inferred heartbeat:

- The core the inferred heartbeat was found on.
- **Period:** The period at which the function was found to execute.
- **Variance:** How close the function was to a regular period.
- **Thread:** The ID of the thread which the function was running on.
- The name of the function.

11.7 Application Information

The **Application Information** provides general data about the running application. The following items are recorded:

- **Average Frame Rate:** The average frame rate seen while profiling. To get a value here, frame marking must be enabled.
- **Application Build SDK:** The SDK that was used to build the application that was profiled.

11.8 Profiler Information

The Profiler Statistics section shows information about the profiler runtime that was used to record the data in the profile. While not directly useful most of the time, these values are useful when reporting problems seen with the profiler.

- **Profiler Version:** The version of the profiler used to record the data.
- **Profiler Build Date/Time:** The date and time that this version of the profiler library was built.
- **Profiler Build SDK:** The SDK used to build the profiler.

11.9 Notes

The **Notes** text box at the bottom of the **Info** tab allows you to store user notes with a profile. These notes are saved along with the profile when it is saved with the **Save** button.

If there are no notes, the **Notes** section will be collapsed with only the header showing. To expand the **Notes** section, either click the up arrow on the right side of the **Notes** header or drag the **Notes** header upward with the mouse.

12 Checkers Tab

The **Checkers** tab contains a list of checkers that have been run against the profile to detect potential problems.

Whenever a new profile is received from the dev kit or a profile is loaded, the checkers will be automatically run against the profile. As checkers complete, any checkers that found a problem show their results in the tab. Checkers that did not find any issues are not displayed.

12.1 Checkers Tab Toolbar

The following controls are on the **Checkers** tab toolbar.

12.1.1 Manage Checkers

Clicking on the **Manage Checkers** button will display all checkers in the **Checkers Tab**, including those disabled and those that did not have any results. This allows disabled checkers to be re-enabled as well as the ability to modify checker parameters for checkers that did not have any results.

When in this management mode, the state of each checker is displayed on the far right side of the checker title. If the checker would normally display a result on a standard run, no special label is shown on the right side. If the checker is Enabled but did not find anything, the text **No Results** is displayed. If the checker is Disabled the text **Checker Disabled** is displayed.



Note: A Disabled checker does not run while in management mode. To see potential results for the checker, it must be Enabled.

12.1.2 Status

To the right of the **Manage Checkers** button is a status area that will display the current status of the running checkers.

There are two different status types that might be displayed:

- `Waiting for ELF processing to finish`: This message indicates that there is still some background processing occurring on the ELF. This is most commonly attributed to the static analysis of the game, looking for static function calls in each function's code.
- `Running checker checker_name`: This message indicates which checker is currently being run on the profile data.

12.2 Right-Click Context Menu

A right-click in the tab will bring up a context menu containing a variety of options.

The context menu for each checker is generated depending on what kind of values are stored in the checker's results table. Any of the following items may be included in the context menu for an individual checker.

A context menu item will be disabled if the specific item in the checker results does not support the operation.

12.2.1 Select Function

The **Select Function** option results in the clicked function becoming selected in the **Functions** tab.

12.2.2 Show Details

The **Show Details** option will open a new [Function Details Window](#), showing the details of the clicked function.

12.2.3 Show Assembly

The **Show Assembly** option will move focus to the [Assembly Tab](#), showing the assembly of the clicked function.

12.2.4 Load per Frame Analysis

The **Load per Frame Analysis** option will select all relevant functions for this table entry and show the load per frame for the entire group in the **Sample Graph**. This will also cause the **Sample Graph** tab to become visible.

To reset the view, click the Reset button on the [Analysis Tab](#).

12.3 Severity

The severity of each checker can be determined by its icon and its position in the list. The most severe issues appear toward the top of the list.

- Items with the purple chevron are the most severe items and are serious issues.
- Items with the red square are severe items and could indicate serious issues.
- Items with the orange diamond are major issues and should be investigated further.
- Items with the yellow triangle represent minor issues or issues that are not known to be serious.
- Items with the green circle are presented for informational purposes only.

12.4 Checker Results

When a checker finds items to report, an entry is added into the **Checkers** tab. Each entry contains a specific set of items.

12.4.1 Rerun Checker Button

The **Rerun Checker** button forces the checker to rerun with the current settings. This process will remove the existing checker results from the list, obtain new results, add the new results to the list and show you those results.

12.4.2 Copy Text Button

Clicking the **Copy Text** button will copy the **Summary**, **Explanation**, and **Suggestions** to the clipboard.

12.4.3 Copy Table Button

Clicking the **Copy Table** button will copy the **Results Table** to the clipboard. The copied text contains tab-separated values and can be paste directly into spreadsheet applications.

12.4.4 Disable Checker Button

The **Disable Checker** button keeps the checker from running and displaying in the **Checker Tab**. The Enabled/Disabled state of the checker is saved in the settings and will persist between profiler runs.

To re-enable the checker, use the [Manage Checkers](#) button in the **Checker Tab Toolbar**.



Note: Enabling or Disabling a checker will result in that checker being automatically rerun.

12.4.5 Summary

All checkers contain a brief **Summary** of the results that were found. If no results were found, this will be stated in the **Summary**.

If a checker takes longer than one second to run, the time will be reported at the bottom of the **Summary**.

12.4.6 Configurable Parameters

When a checker has settings that can be configured, spinner controls will appear beneath the **Summary**. Each spinner starts at the default value and has a maximum and minimum allowed value. Adjusting the parameter can result in better checker results for a specific application. Any changes to these parameters are saved for use in subsequent runs of the profiler. To reset the parameters to the default values (and overwrite any saved values), click the **Reset to Defaults** button.

12.4.7 Reset to Defaults Button

The **Reset to Defaults** button is a quick way to set all of the Checker [Configurable Parameters](#) back to their default values.

12.4.8 Results Table

Each checker creates its own **Results Table**. There are no set columns that are guaranteed to appear in a checker.

Columns that commonly appear in checkers include:

- **Self:** The amount of time attributed to this function in the profile results.
- **Function:** The name of the function found.

12.4.8.1 Sorting the Results Table

Each **Results Table** comes pre-sorted using a default sort determined to be the most generally useful for the checker results.

If the default sorting is insufficient, or to get a different view the results, it is possible to change the sorting used in the results. Click on any of the headers in the table to sort the results by that column. When this action is performed an arrow will appear in the column to indicate the sorted column and the sort direction. Clicking on the same column again will result in the sort direction changing.

12.4.9 Explanation

The details in the **Explanation** describe what the checker is looking for and why it is looking for it.

12.4.10 Suggestions

The details in the **Suggestions** provide guidance on how to change your application to correct the issues reported by the checker.

13 Call Tree Tab

This tab shows the discovered call tree based on the profiled sample data. This tab only appears if the profile was taken with **Callstack** data turned on (see [Callstack Toggle Button](#)).

13.1 Call Tree Tab Toolbar

The following controls are on the **Call Tree** tab toolbar.

13.1.1 Prefix Selection

The prefix before each function name is determined by the three **Prefix** toggle buttons in the **Call Tree** toolbar. By default, the **Total** and **Self** prefixes are turned on. To toggle a prefix on and off, click on the prefix button. If all buttons are set to off, just the function name will be shown. The prefix choices are:

Total: This shows the time spent in each function and the functions it calls.

Self: This shows the time spent in each function, not including the functions it calls.

Sub: This shows the time spent in functions that are called by each function.

Note that the functions at each level of the tree are sorted based on the prefix. If **Total**, **Self**, or **Sub** are selected, the sorting is from highest to lowest of the leftmost prefix. If no prefix is selected, the sorting is alphabetical, based on the function name.

13.1.2 Expand All Button

The **Expand All** button expands all possible items in the view.

13.1.3 Collapse All Button

The **Collapse All** button collapses all possible items in the view.

13.1.4 Auto Expanding

By default, the **Call Tree** will not auto-expand, however you can turn on this behavior with the **Auto Expand** drop-down box. If auto expanding is turned on, the tree will expand and/or collapse based on which functions are selected with mouse clicks or highlighted as a result of the **Find** functionality (in either the **Functions** tab or **Call Tree** tab).

The **Auto Expand** drop down box has three choices:

- **Auto Expand Off:** The **Call Tree** will not expand or collapse based on functions being selected/deselected or highlighted by **Find**.
- **Auto Expand Only:** The **Call Tree** will expand (but not collapse) based on functions being selected or highlighted by **Find**.
- **Auto Expand/Collapse:** The **Call Tree** will expand and collapse based on functions being selected/deselected or highlighted by **Find**.

Additionally, you can quickly expand and collapse the entire tree with the **Expand All** and **Collapse All** buttons.

13.1.5 Copy Button

The **Copy** button will copy to the clipboard whatever is currently shown using the current sorting and filtering. You can then paste the content into an email, document, or Excel spreadsheet. Because the

copied content uses tab delimiters, it will retain the column organization when pasted into an Excel spreadsheet.

The **Copy** button takes into account what is currently selected. If any items are selected at the time the **Copy** button is pressed, then only the selected items will be copied. If nothing is currently selected, everything will be copied.

13.1.6 Copy All Button

The **Copy All** button will copy to the clipboard whatever is currently shown using the current sorting and filtering. You can then paste the content into an email, document, or Excel spreadsheet. Because the copied content uses tab delimiters, it will retain the column organization when pasted into an Excel spreadsheet.

Unlike the **Copy** button, **Copy All** does not respect selected items. This allows for one to copy the entire visible call tree at any time.

13.2 Finding Functions

Type any string into the text box and a case-insensitive search is performed for the items that contain the string. Only items matching the search text will be displayed. To negate the results (show everything that does not match the filter), click on the exclamation point button to the right of the text box. If there are no matches for the search, the text box will be highlighted orange.

Regular expressions are also supported (using .NET RegEx, which is Perl 5 compatible). To have the filter interpreted as a regular expression, click on the **RegEx** button to the right of the text box. Unlike basic filters, regular expressions are case-sensitive.

Helpful function filters can be found in the [Helpful Simple Filters](#) and [Helpful Regular Expression Filters](#) tables. Please see the [Quick Regular Expression Primer](#) for more information on regular expressions.



Note: The **Find** operation here is different than the **Filter** operations on the **Functions** and **Code Coverage** tabs. In the **Call Tree** view all functions are always visible. Instead, only the found results of the search are highlighted.

13.2.1 Saved Filters

Saved filters allow for using the same filters between profiling sessions, with any saved filters being loaded when the profiler is opened. By default, the **Filter Drop Down** box is loaded with example filters that you can customize. If you want to save your own filter in the drop down box, press the **Save** ("+") button after typing the filter (the **Negate** and **RegEx** button states will also be saved with the filter text). To delete a filter from the drop down box, select it first, then press the **Delete** ("-") button.



Note: The saved filters in the **Functions** tab will also be available in the **Call Tree** and **Code Coverage** tabs. However, please note that the **Filter Drop Down** box in the **Threads** tab and **Counters** tab are not shared since they are used to search for thread and counter names, not function names.

13.3 Function Selection

Clicking on a function in the **Call Tree** will highlight the function, similarly to how function selection works in the **Functions** tab (up to 45 functions can be highlighted). All instances of the highlighted function will be highlighted in the **Call Tree**. If auto-expanding is on, then the tree will be expanded to show all instances of the selected function.

13.4 Right-Click Context Menu

A right-click in the tab will bring up a context menu containing a variety of options.

13.4.1 Manual Expanding

The call tree works like a typical Windows tree control, where the tree can be expanded or collapsed at any node by clicking the plus or minus sign. Additionally, you can right click on a single node and choose to expand or collapse downward in the tree from that node.

13.4.2 Show Assembly

The **Show Assembly** option will move focus to the [Assembly Tab](#), showing the assembly of the clicked function.

13.4.3 Show Details

The **Show Details** option will open a new [Function Details Window](#), showing the details of the clicked function.

13.5 Partial Call Tree

When a function is selected, it will also appear in the **Partial Call Tree** at the bottom of the **Call Tree Tab**. There are three different views available in this sub-window which allow for differing ways to view the data.

The **Partial Call Tree** window can be resized by dragging the header up and down. It can also be collapsed by clicking on the down arrow on the right side of the header and similarly expanded by clicking on the up arrow.

Forward from Selected

This view treats the selected function as the root and expanding it successively allows you to travel down the call tree. This produces a similar view to that of the **Call Tree**. The difference, though, is that if there are multiple calling functions in the **Call Tree**, this view will combine the selected function into a single entry. This may allow for easier viewing of the calls made by the selected function.

The **Total**, **Self**, and **Sub** values in the view are calculated as the percentage of the profile where the specific tree view was present.

Reverse from Selected

This view treats the selected function as the root and expanding it successively allows you to travel up the call tree. However, this is not just a simple reversal of the call tree, but instead separates and groups functions based on how they are called in order to show relationships that would otherwise be hard to see. For example, a vector normalize function might appear 10 times in the **Call Tree**, but it might be called from only 3 different functions. In this case, the vector normalize function will be listed once, with a child node for calling function.

The **Total**, **Self**, and **Sub** values in the view are calculated as the percentage of the profile where the specific tree view was present.

Legacy Reverse

Similar to **Reverse from Selected**, only in this case, the vector normalize function will be listed 3 times, with each entry showing the total time called from each parent function. This is the legacy behavior from previous versions of the profiler.

The **Total**, **Self**, and **Sub** values in this view are calculated as the percentage of the whole profile.

13.6 Units

To change the units displayed before each item, use the **Unit** buttons in the **Home** tab of the ribbon bar (see [Units](#)).

14 Sample Graph Tab

The **Sample Graph** graphically depicts selected functions, threads, counters, and instrumented functions/code blocks. Click on each item within their respective tabs to draw them in the **Sample Graph**.

When **Sample by Time** is selected, functions and threads each will be depicted as a rectangle starting at the time it was sampled and ending at the following sample. As sampling time can vary, the widths of these rectangles may vary, too. Additionally, if the same function or thread is sampled multiple times in a row, it will appear as an elongated rectangle.

If one of the **Sample by Performance Counter** choices was used from the **Sampling Strategy** drop down box, then function and thread samples will be drawn as a single vertical line at the time it was recorded.

Sample by Performance Counter sampling may record samples at an irregular rate with respect to time (with very small or very large gaps between some samples). This is expected behavior since performance counter events occur irregularly.

A **Load per Frame** graph will also be drawn for each selected function and thread if a logged heartbeat or inferred heartbeat is selected for that core. The **Load per Frame** allows you to quickly assess the character of a given function on a frame-by-frame basis. The dark line shows the actual data and the highlight shows the margin of error (which adaptively follows the line to give an indication of the data's average value). Additionally, the selected **Units** in the **Home** tab of the **Ribbon Bar** will control the vertical axis of the **Load per Frame** graph.



Note: The vertical scale of the **Load per Frame** graph can be controlled by a drop down selection in the **Sample Graph** toolbar.

For counters, a line graph will be drawn between each value recorded. Most counters listed under [Types of Counters](#) accumulate between samples and are cleared after the sample is recorded, so the value of the counter at the time of the sample is only correlated with the function sampled. Conversely, user logged counters from inside the code are recorded instantaneously and graphed as such.

For instrumented functions and code blocks, the graphs depend on the type of data selected. However, since the recorded data is a constant value between the start and end a function or code block, each recorded chunk is drawn as a rectangle at the height of the value. The left edge of the rectangle is the time the function or code block was entered and the right edge of the rectangle is the time the function or code block was exited (thus making the width of the rectangle equal to the duration of the of the function call or code block). For example, when graphing **Time per Call** for an instrumented function, the height of the rectangle is the time spent in that function call (and all functions it calls) and the width of the rectangle is also the time spent in that function call. When graphing **Function Calls** for an instrumented function, the height of the rectangle is always 1 (because it was called one time for a given function call) and the width is the duration. When graphing a performance counter the height of the rectangle is equal to the counter value for the function call or code block and the width is the duration.

The **Sample Graph** tab is drawn using DirectX 10. If there is a problem initializing DirectX within the **Sample Graph** tab, an error message will be output to the screen containing details of the problem. Some DirectX troubleshooting is covered in [Troubleshooting](#).

14.1 Sample Graph Tab Toolbar

The following controls are in the **Sample Graph** tab toolbar.

14.1.1 Mouse Mode

The left section in the **Sample Graph** toolbar controls the mouse mode within the **Sample Graph**. There are four modes:

1. **Scroll Graph:** In this mode, left-click-and-drag will scroll the graph up, down, left, and right.

Keyboard modifiers:

- Holding the Ctrl key will temporarily engage the **Measure Time** mouse mode until Ctrl is released.

2. **Select Frames:** In this mode, left-click or left-click-and-drag will select heartbeat frames. To deselect frames, left-click or left-click-and-drag on an already selected frame.

Keyboard modifiers:

- Frames can also be deselected by holding the Alt key along with a left-click-and-drag.
- Holding the Ctrl key will temporarily engage the **Scroll Graph** mouse mode until Ctrl is released.

3. **Select Time:** In this mode, left-click-and-drag to select regions of time. Note that times will be snapped to sample boundaries within the core that was initially clicked on. To deselect time, left click or left click and drag on an already selected time.

Keyboard modifiers:

- Time can also be deselected by holding the Alt key along with a left click and drag.
- Holding the Ctrl key will temporarily engage the **Scroll Graph** mouse mode until Ctrl is released.

4. **Measure Time:** In this mode, left click and drag to measure regions of time. When the left mouse button is released, the region will disappear.

Keyboard modifiers:

- Holding the Ctrl key will temporarily engage the **Scroll Graph** mouse mode until Ctrl is released.

In the **Select Frames** and **Select Time** modes, after selecting an area the **Functions**, **Threads**, **Call Tree**, and **Counters** tabs will recalculate based off of this selection.



Note: Some controls work in any mouse mode.

- The left and right arrow keys can be used to page left or right and the up and down arrow keys to zoom in and out.
- A function can be double-clicked to expand its direct children. This option is also available in the right-click context menu.



Caution: When selecting individual frames or times, the amount of statistical data is greatly reduced and the margin of error will dramatically rise. Double check the margin of error by enabling the [Error button](#) on the **Home** tab of the ribbon bar.

14.1.2 Deselect All Button

To quickly deselect all selected frames or selected periods of time, click on the **Deselect All** button.

14.1.3 Icicle Graph

The **Icicle Graph** is a visualization of the call stack at every recorded sample. To best view the graph, select at least one function on a core, zoom into a single frame or set of frames, and check the **Icicle Graph** toggle button. It is recommended to zoom into a single frame or small set of frames since the graph is extremely detailed. The name **Icicle Graph** comes from the look of the graph, with root functions at the top and callstack leaves toward the bottom at the tips.

As a visualization of the callstack, the **Icicle Graph** shows the calling behavior from the root function down to the leaf of each callstack. To see individual function names and statistics (depth, total, and self), use the mouse to hover over the segment you are interested in. The color of each function is based on the automatic system classification performed by the profiler, with the color key located above each core header label. The automatic system classification might be incorrect for some functions, but on average it performs well and is helpful in understanding which systems are executing.

Icicle Graph rendering for multiple cores can be slow and will make the profiler feel sluggish to zoom or scroll. To keep rendering fast, it is recommended that you only draw an **Icicle Graph** for one core at a time

(by not selecting functions/threads/counters on multiple cores). Since the **Icicle Graph** only draws on cores where a function, thread, or counter are selected, just restrict selection to one core.

When looking at individual frames, it is often the case that the **Load per Frame** graph is not helpful at this zoom level. To quickly hide the **Load per Frame** graph, deselect the **Load Graph** button in the **Sample Graph** toolbar.

14.1.4 Icicle Selected Graph

The **Icicle Selected** graph is similar to the **Icicle Graph**. However, instead of displaying colors for every function, only the currently selected functions are colored in the **Icicle Graph**. This has the benefit of making it easy to see where in a callstack a specific function was called at a particular time.

14.1.5 Stack Graph

The **Stack Graph** provides a quick and easy way to track the depth of the stack across a profile. The color of each bar is based on the leaf function's classification, with the color key located above each core header label.

If the profiler was unable to obtain the stack depth, the graph values will be fixed at 0.

14.1.6 Load Graph

This button provides a quick way to hide or show the **Load per Frame Graph** in the **Sample Graph**. For more in-depth control of the **Load per Frame Graph** please see [Samples, Load, and Counters Scale](#).

14.1.7 Heatmap Mode

Heatmap Mode stacks all frames on top of each other, drawing each sample faintly so that repetitive trends can be visualized. This mode is most helpful when sampling at a low rate (below 1000 samples per frame), since within a single frame there won't be enough data to see the true behavior. However, by stacking the samples on top of each other, you can leverage dozens or hundreds of frames to build up a visualization of how your game performs within a frame.

Since **Heatmap Mode** stacks frames, cores that do not have a heartbeat frame selected will be hidden in this mode. To force a heartbeat framing across all cores, go to the **Heartbeats** tab in the ribbon bar, find a desirable heartbeat on any core, then click **Apply to All Cores**.

The vertical black line that appears in **Heatmap Mode** represents the end of each stacked frame. The line is drawn faintly so that many stacked frames will accumulate to make it darker as they overlap each other. This is a good way to judge if frames are generally taking the same amount of time (by stacking directly on top of each other), whether they fall into several clusters of frame rates, or whether they are wildly scattered.

Often a profile contains a mix of frame rates which can make the **Heatmap Mode** difficult to understand. This is where you should leverage the **Group Selection** options in the **Heartbeats** tab. By selecting only singular frame rates (like **20Hz**), you can stack only these slow frames, thus eliminating noise from irrelevant frames. Additionally, this is a really good time to also use the **Story Analysis** in the **Analysis** tab only select important functions within these particular frames.

In **Heatmap Mode**, function names can get in the way of seeing the data. One option is to select **Short Names** or **Hide Names** in the **Sample Graph** ribbon bar.

14.1.8 Self Samples / Total Samples / Mixed Samples

This drop down selection only appears for **Callstack** profiles that contain callstack data. With this drop down you can select between graphing samples by **Self Samples**, **Total Samples**, or **Mixed Samples**.

- Functions graphed as **Self Samples** will show a sample marker each time the function was sampled as a leaf function (when the program was stopped and sampled, execution was in that function).

- Functions graphed as **Total Samples** will show a sample marker each time the function was sampled as part of a callstack, regardless of whether it was a leaf function.
- Functions graphed as **Mixed Samples** will show the same samples as **Total Samples**, but function leaf samples (**Self Samples**) are drawn in a darker color.

For example, the function `main` might have very few sample markers graphed when viewed with **Self Samples**, but it might have a very large number of sample markers graphed when viewed with **Total Samples** or **Mixed Samples**.

14.1.9 Adjusting Name Widths

When a function or thread is graphed, its name appears on the left side of the **Sample Graph**. The **Names** drop-down in the **Sample Graph** toolbar contains four settings that control how function name or thread name are displayed:

- **Hide Names**: This setting will hide all function names and thread names by not drawing them.
- **Short Names**: This setting will display shortened function names (limiting each to 100 pixels wide). For thread names, only the thread's ID will be shown.
- **Long Names**: This setting will display the function name without arguments. Thread names will be shown in full (ID and name).
- **Full Names**: This setting will display the entire function name including arguments. Thread names will be shown in full (ID and name).

14.1.10 Combined Graphs

The **Combined Graphs** allow you to view all functions or threads for a given core drawn on the same horizontal row, which may help with understanding the execution sequence and for spotting gaps as a result of unselected functions or threads.

The **Combined Graphs** drop down box allows you to select the drawing behavior of combined function or thread graphs. The drop down contains three choices, with **Hide Combined** as the default selection:

- **Hide Combined**: No combined graph is drawn.
- **Show Combined**: Combined graphs are drawn for functions and threads.
- **Combined Only**: Combined graphs are drawn for functions and threads, while individual graphs for functions and threads are not drawn.

14.1.11 Graph Tool Tips

The **Graph Tool Tips** drop down allows you to specify the type of tool tips you want to see when hovering over a function row, thread row, or **Iceberg Graph** function.

- **Hide Tool Tips**: No tool tips are shown.
- **Show Name Only**: Only the name is shown.
- **Show Brief Stats**: The name, sample statistics, and frame number are shown.
- **Show Full Stats**: The name, thread, system classification, spike detection, periodicity, sample statistics, frame number, and exact time within the profile are shown.
- **Show Fn at Time**: The function or callstack will be shown at the exact time pointed at. This ignores the function or thread row and instead focuses on the horizontal position of the mouse within the duration of the profile.

14.1.12 Line Highlight

The **Line Highlight** drop down allows you to specify the type of highlight that appears in the **Load per Frame** graph.

- **Hide Line Highlight**: No highlight is drawn.
- **Margin of Error**: The highlight represents the margin of error around the data.

- **Min/Max Value:** The highlight is a horizontal band representing the min and max values of the data.

14.1.13 Sample Gaps

The **Sample Gaps** drop down allows you to show or hide sample gaps in the profile. Sample gaps are determined by looking at when samples were recorded by the profiler and seeing when the gap between the samples significantly exceeds expected values. One possible cause of this behavior is if interrupts are disabled.

- **Hide Sample Gaps:** Sample gaps are not drawn.
- **Show Sample Gaps:** Sample gaps are drawn.

14.1.14 Samples, Load, and Counters Scale

The vertical size or scale of graphs—**Samples**, **Load**, or **Counters**—can each be set using its **Vertical Scale** drop-down in the **Sample Graph** toolbar. All changes are saved in the settings for the next time the profiler is loaded. There is also an option to **Hide** all line graphs within each **Vertical Scale** drop-down.

A fast way to hide/show each is to click on the label names themselves (**Samples**, **Load**, or **Counters**). Clicking on the label name again will reset the scale back to its previous value.

Practical vertical scales will be in the range of **1x** to **5x**. Use the larger vertical scales to precisely measure a particular value, since the y-axis will contain more resolution and unit labels.

14.1.15 View Control

The **Sample Graph** toolbar contains buttons to control the view:

- **Hard Scroll Left (left pointing arrow):** Scrolls the view all the way to the left.
- **Scale to Fit (left and right arrows):** Shows the entire profile, scaled horizontally to fit the window.
- **Hard Scroll Right (right pointing arrow):** Scrolls the view all the way to the right.

14.2 Timeline

At the bottom of the **Sample Graph** is the **Timeline**. This shows the entire zoomed out profile as a frame rate graph. Overlaid on the **Timeline** is a window representing your current zoomed in area of the graph above. You can click and drag the window on the timeline as if it was a scrollbar. Also, clicking once on either side of the window will move it one window width to the left or right. The mouse wheel also works within the **Timeline** area to zoom in and out.

If the frame rate graph is not being displayed in the **Timeline**, it is because the game is not marked for frames properly. See [Frame Marking on Wii U](#) for more information on how to do this. Alternatively, you can go to the **Heartbeats** tab and for a given heartbeat on a given core, you can choose **Use as Frame Rate** to have that heartbeat be drawn as the frame rate in the **Timeline**.

14.3 Zoom Control

The zoom (horizontal scale) can be controlled using the mouse wheel. The current level of zoom is displayed in the **Sample Graph** toolbar. If you want to zoom all the way out, the easiest method is to click the **Scale to Fit** button in the **Sample Graph** toolbar.

An alternative control for the zoom is the up and down arrow keys. This can be a good way to quickly zoom into the graph.

14.4 Right-Click Context Menu

A right-click in the tab will bring up a context menu containing a variety of options.

When the menu is up, there will be a highlight placed on the current function (orange horizontal rectangle) and the current time (orange vertical line). Particular choices in the context menu will affect either the current function or current time, so these are marked for reference.

14.4.1 Expand Direct Parents

The **Expand Direct Parents** option will select the parents of the currently selected function and place them above the currently selected function. Since a function can be called from many different parents, all parents are selected.

14.4.2 Expand Direct Children

The **Expand Direct Children** option will select the children of the currently selected function and place them below the currently selected function. A quicker way to accomplish this is to simply double click on the any function in the **Sample Graph**.

14.4.3 Deselect Direct Children

The **Deselect Direct Children** option is like an undo for the **Expand Direct Children** option.

14.4.4 Reorder

There are several **Reorder** options:

- **Chronologically**: Reorders all selected functions based on their first occurrence within the frame that right clicked. This is the same ordering used in the **Story Analysis**.
- **By Magnitude**: Reorders all selected functions based on how often each function was sampled. This is similar to the order in the **Functions** tab when it is sorted by **Total**.
- **By Periodicity**: Reorders all selected functions based on how periodic they are. Functions that are very periodic will be placed toward the top and functions that are sampled very irregularly will be placed toward the bottom.
- **By Function Name**: Reorders all selected functions alphabetically by function name. This might help to locate a particular function if you know the name and there are many functions displayed in the **Sample Graph**.

14.4.5 Deselect

There are several **Deselect** options:

- **This Function**: Deselects the currently highlighted function.
- **All Other Functions**: Deselects all functions other than the currently highlighted function.
- **All Functions On This Core**: Deselects all functions on the core of the currently highlighted function.
- **All Functions On Other Cores**: Deselects all functions on other cores of the currently highlighted function.
- **Above This Function**: Deselects all functions above the currently highlighted function.
- **Below This Function**: Deselects all functions below the currently highlighted function.

14.4.6 Move

There are several **Move** options:

- **Move Function Up**: Moves the currently highlighted function up one row.
- **Move Function Down**: Moves the currently highlighted function down one row.

14.4.7 Show Assembly

The **Show Assembly** option will move focus to the [Assembly Tab](#), showing the assembly of the clicked function.

14.4.8 Show Details

The **Show Details** option will open a new [Function Details Window](#), showing the details of the clicked function.

14.4.9 This Core At This Time

There are several **This Core At This Time** options:

- **Find Sampled Leaf Function:** Selects the sampled leaf function on this core at the currently highlighted time. The function will be placed at the top of the currently shown functions.
- **Find All Sampled Functions:** Selects all sampled functions (the callstack) on this core at the currently highlighted time. The functions will be placed at the top of the currently shown functions.

14.4.10 All Cores At This Time

There are several **All Cores At This Time** options:

- **Find Sampled Leaf Functions:** Selects all sampled leaf functions on all cores at the currently highlighted time. The functions will be placed at the top of the currently shown functions.
- **Find All Sampled Functions:** Selects all sampled functions (the callstack) on all cores at the currently highlighted time. The functions will be placed at the top of the currently shown functions.

14.4.11 Select Only This Heartbeat Frame

The **Select Only This Heartbeat Frame** option will deselect all frames/times and then select the current frame that is being hovered over with the mouse. This can be useful when you want to quickly focus on a specific frame for **Story Analysis** or **Spike Detection** (select the frame first and then activate the analysis).

If you want to select more than just one frame, change the mouse mode to **Select Frames** in the **Sample Graph** toolbar and use the left mouse button to select or deselect frames.

15 Code Coverage Tab

The **Code Coverage** tab allows you to track which functions have been seen and which have not been seen in a given profile. This can be helpful when trying to determine coverage for testing or for discovering unanticipated behavior.

Please note that loading a saved profile will clear the lists and they will be recalculated based on the saved profile. The lists are not saved with a profile but are instead generated based on the profile.

15.1 Functions Seen During Profiling

This window contains the names of functions that have been seen during profiling. The number of functions in the list is reported in the heading along with the combined code size. The combined code size is only visible when there is no filter applied.

15.1.1 Copy Button

The **Copy** button will copy to the clipboard whatever is currently shown using the current sorting and filtering. You can then paste the content into an email, document, or Excel spreadsheet. Because the copied content uses tab delimiters, it will retain the column organization when pasted into an Excel spreadsheet.

The **Copy** button takes into account what is currently selected. If any items are selected at the time the **Copy** button is pressed, then only the selected items will be copied. If nothing is currently selected, everything will be copied.

15.2 Functions Not Seen During Profiling

This window contains the names of functions that have not been seen during profiling. The number of functions in the list is reported in the heading along with the combined code size. The combined code size is only visible when there is no filter applied.

15.2.1 Copy Button

The **Copy** button will copy to the clipboard whatever is currently shown using the current sorting and filtering. You can then paste the content into an email, document, or Excel spreadsheet. Because the copied content uses tab delimiters, it will retain the column organization when pasted into an Excel spreadsheet.

The **Copy** button takes into account what is currently selected. If any items are selected at the time the **Copy** button is pressed, then only the selected items will be copied. If nothing is currently selected, everything will be copied.

15.3 Filtering Functions

Type any string into the text box and a case-insensitive search is performed for the items that contain the string. Only items matching the search text will be displayed. To negate the results (show everything that does not match the filter), click on the exclamation point button to the right of the text box. If there are no matches for the search, the text box will be highlighted orange.

Regular expressions are also supported (using .NET RegEx, which is Perl 5 compatible). To have the filter interpreted as a regular expression, click on the **RegEx** button to the right of the text box. Unlike basic filters, regular expressions are case-sensitive.

Helpful function filters can be found in the [Helpful Simple Filters](#) and [Helpful Regular Expression Filters](#) tables. Please see the [Quick Regular Expression Primer](#) for more information on regular expressions.

15.3.1 Saved Filters

Saved filters allow for using the same filters between profiling sessions, with any saved filters being loaded when the profiler is opened. By default, the **Filter Drop Down** box is loaded with example filters that you can customize. If you want to save your own filter in the drop down box, press the **Save** ("+") button after typing the filter (the **Negate** and **RegEx** button states will also be saved with the filter text). To delete a filter from the drop down box, select it first, then press the **Delete** ("-") button.



Note: The saved filters in the **Functions** tab will also be available in the **Call Tree** and **Code Coverage** tabs. However, please note that the **Filter Drop Down** box in the **Threads** tab and **Counters** tab are not shared since they are used to search for thread and counter names, not function names.

15.4 Right-Click Context Menu

A right-click in the tab will bring up a context menu containing a variety of options.

15.4.1 Copy Function Name to Clipboard

This will copy the full function name with arguments to the clipboard.

15.4.2 Select Function

The **Select Function** option results in the clicked function becoming selected in the **Functions** tab.

15.4.3 Show Assembly

The **Show Assembly** option will move focus to the [Assembly Tab](#), showing the assembly of the clicked function.

15.4.4 Show Details

The **Show Details** option will open a new [Function Details Window](#), showing the details of the clicked function.

16 Assembly Tab

The **Assembly** tab allows you to see where in a function samples occurred. This can be helpful in tracking down specific problem spots within a function.

16.1 Assembly Tab Toolbar

The following controls are in the **Assembly** tab toolbar.

16.1.1 Copy Button

The **Copy** button will copy to the clipboard whatever is currently shown using the current sorting and filtering. You can then paste the content into an email, document, or Excel spreadsheet. Because the copied content uses tab delimiters, it will retain the column organization when pasted into an Excel spreadsheet.

The **Copy** button takes into account what is currently selected. If any items are selected at the time the **Copy** button is pressed, then only the selected items will be copied. If nothing is currently selected, everything will be copied.

16.1.2 View Controls

In order to more readily facilitate looking through large functions with few samples, a **Previous** and **Next** button are provided. These will scroll to the previous or next off-screen line of assembly, respectively, that has non-zero count in either self or sub.

16.1.3 Show Source

It is possible to view interleaved source code with the assembly. This information is obtained from the ELF files stored on your hard drive. As such, there are a few limitations on viewing source code.

This button works as a toggle. On each profiler startup it will default to being disabled. If the button is enabled, selecting a new function for viewing will automatically attempt to pull in source code for the function.



Note: Clicking on **Show Source** will force the sorting to **Address** and then interleave the source code.

16.1.3.1 Limitations on Viewing Interleaved Source Code

Provided here is a list of limitations on viewing interleaved source code.

- The Green Hills MULTI tools must be installed on the computer. The profiler makes use of the `gaddr2line` tool in the MULTI install path to obtain which source code line pertains to each assembly instruction. If the MULTI tools cannot be located, the message "Unable to locate MULTI tools" will be displayed.
- The original ELF files must exist on your computer in the same locations in which they were previously found when the profile was initially taken. If the required ELF cannot be located, the message "Unable to locate ElfName" will be displayed.
- The ELF on the computer must have a date that is older than the date/time at which the profile was taken. If the date on the ELF is newer, then it is not possible that the ELF corresponds with the current profile data. If this situation occurs, the message "ELF modified date/time is newer than the profile date/time" will be displayed.

- Individual source files are not checked for date/time stamps. As such, it is possible that the source file has been modified and line numbers changed since the ELF was created. In this case, incorrect lines will be displayed in the assembly view. Be mindful of this limitation while using the interleaved source code view.

16.2 Selecting a Function to View

When the **Assembly** tab is open, click on a function in the **Functions** tab to show its data in the **Assembly** tab. The currently selected function is shown to the right of the text **Selected Function** in the toolbar area of the **Assembly** tab. Please note that the selected function will not be highlighted in the **Functions** tab.

Functions that are RPL internal functions are not viewable. If one of these functions is selected, a message will display in the assembly window along with **Self** and **Sub** counts for the function. These should match the counts visible in the **Functions** tab.

16.3 Assembly View

This is the area of the tab which actually displays the assembly source code. Each line of source code contains the following information:

- **Address**: The address at which the line of assembly appears in memory.
- **Self**: The number of samples taken when the program counter was at that address, executing that line of assembly.
- **Sub**: The number of samples taken where this function was above the sampled instruction in the **Call Tree**. As such, this number will only appear on branch-with-linking instructions in **Callstack** profiles.
- **Assembly**: The assembly instruction that is stored at that address.

The Wii U CPU Profiler uses a different disassembler than the one provided by the provided compiler. As such, a direct comparison of the assembly against the compiler tools version will result in slightly different text. However, the difference should be in the mnemonic that is used (or not used as the case may be).

The columns can be resized and reordered as much as desired. However, when the **Copy** function is used, the default column ordering and sizes are used.

16.3.1 Highlighting

In order to make it easier to visually assess which areas of the assembly are seen more often in profiling, the line with the most **Self** samples is colored the darkest red color. Lighter shades are used for lessening **Self** hit counts. If there are more **Sub** samples than there are **Self** samples, then the color will change varying shades of a yellow-green color.

16.3.2 Sorting

Rudimentary sorting is available for the assembly view. The default sorting is by ascending **Address**. Any time a new function is selected to view the sorting mode is set back to the default.

To change the sorting, click on the column header to sort by that column. Each column can only be sorted a single direction. **Address** and **Assembly** both use an ascending sort while **Self** and **Sub** both use a descending sort.



Note: When any of the column headers are clicked **Show Source** is automatically disabled.

16.4 Right-Click Context Menu

A right-click in the tab will bring up a context menu containing a variety of options.

Options that target a specific function are only enabled when the assembly instruction is a branch.

Options that target data are only enabled when there is a literal pool in the function.

16.4.1 Select Target Function

The **Select Target Function** option results in the clicked function becoming selected in the **Functions** tab.

This option is only enabled if the target function was seen in the profile. If the target function was not seen then the option will be disabled.

16.4.2 Jump to Target

Selecting the **Jump to Target** option will result in the **Assembly View** showing the target address. If the jump occurs within the same function the **Assembly View** scrolls to bring the target address to the top of the window. If the jump is to another function then the **Assembly View** changes the selected function to the target function.

16.4.3 Show Target Details

The **Show Target Details** option will open a new [Function Details Window](#), showing the details of the target function.

16.4.4 Copy Target Function Name

The **Copy Target Function Name** option will copy the full function name with arguments to the clipboard.

16.4.5 View Data As

When a function contains a literal pool the profiler looks at the data in the literal pool and attempts to automatically determine the type of that data. The **View Data As** options allow for overriding the automatically assigned type. The following types are currently allowed for assignment:

- **Int32**: A 32-bit signed integer
- **UInt32**: A 32-bit unsigned integer
- **Single**: A 32-bit floating-point value
- **String**: A 4 byte array of extended ASCII (ISO 8859-1) characters
- **Function**: Treat the data as if it is a function address

17 Console Tab

The **Console** tab can be used to show debug output written from the dev kit. Output is only available when Synced to a dev kit with the **Enable** button toggled to the on position.

The Console tab acts like a console window, with new output appearing at the bottom of the window. Approximately 1000 lines can be buffered before beginning to lose previous lines output.

17.1 Console Tab Toolbar

The following controls are in the **Console** tab toolbar.

17.1.1 Enable Button

By default the **Console** does not display debug output from the dev kit. In order to get debug output, you must toggle the **Enable** button. To stop receiving debug output, toggle the **Enable** button again.

17.1.2 Clear Button

The **Clear** button clears everything from the window.

17.1.3 Copy Button

The **Copy** button will copy to the clipboard whatever is currently shown using the current sorting and filtering. You can then paste the content into an email, document, or Excel spreadsheet. Because the copied content uses tab delimiters, it will retain the column organization when pasted into an Excel spreadsheet.

The **Copy** button takes into account what is currently selected. If any items are selected at the time the **Copy** button is pressed, then only the selected items will be copied. If nothing is currently selected, everything will be copied.

17.1.4 Text Format Selection

The **Text Format** drop-down allows for changing the format used to decode text output from the dev kit. This rule is applied as text is received and is not retroactively applied.

Valid selections are:

- **UTF-8**
- **Shift-JIS**

18 Function Details Window

The **Function Details** window provides additional information about a selected function. The following information is displayed:

- The name of the function.
- The module the function was found in.
- The address of the function.
- The size, in bytes, of the function.
- The number of instructions in the function.
- The functions that this function calls.
- The functions that call this function.

Additionally, if the function was seen during a profile, the following details are also displayed.

- The minimum **Total** % of the function (only displayed for **Callstack** profiles; found by finding highest **Total** % in the call tree for all threads and cores).
- The combined **Self** % of the function (found by summing all **Self** % together across all threads and cores).
- The threads on which the function was seen.

18.1 Selecting a Function to Display

Selecting a function to display in the **Function Details** window occurs by using the right-click context menu. Right-click on the function you wish to view details for and select **Show Details**. The **Functions Details** window will then be displayed. It must be closed before continuing to interact with the rest of the profiler.

The following tabs in the profiler contain a context menu option to show function details:

- [Functions Tab](#)
- [Checkers Tab](#) (some, but not all of the [Results Tables](#))
- [Checkers Tab](#)
- [Call Tree Tab](#) (both regular and reverse call tree)
- [Sample Graph Tab](#)
- [Code Coverage Tab](#)

18.2 Function Details Operations

The following operations are supported from the **Function Details** window.

18.2.1 Show Assembly

The **Show Assembly** option will move focus to the [Assembly Tab](#), showing the assembly of the clicked function.

18.2.2 Select Function

The **Select Function** option results in the clicked function becoming selected in the **Functions** tab.

18.2.3 Traversing Function Calls

To change which function the **Function Details** window is looking at, it is possible to double-click on a function in either the **Calls Functions** or **Called by Functions** lists. This will switch to displaying the details for the selected function. This process can be continually repeated.

19 System Classifications

Every function is automatically classified as to what system it belongs to (Physics, Graphics, Audio, etc). This is primarily accomplished with string matching using commonly accepted terms for each system (for example, "broadphase" for Physics or "matrix" for Math). It is fully expected that some functions will be classified incorrectly, but since it is mostly accurate, it is still immensely helpful in aiding understanding.

System classifications are used in three primary ways:

- **Highlighting Functions:** When **Color Coding** is set to **System**, then all selected functions will be color coded to their system classification when selected.
- **Icicle Graph:** All functions within the **Icicle Graph** are color coded to their system classification.
- **Stack Graph:** All functions within the **Stack Graph** are color coded to their system classification.
- **Quick Filters:** The **Quick Filters** use the system classifications to identify each class of functions.

19.1 Customizing Classifications

How the classifications determine the which functions are part of which system is through the regular expressions in the `bin\QuickFilters\Classifications.xml` file. As our default selections may miss specific parts of some code, we have provided a way to replace the regular expression used. The method is very similar to the one used by [Quick Filters](#).

How to change the regular expression used:

1. Make a copy of the current `Classifications.xml` file and place the copy in the folder denoted by the [Custom Button Directory](#).
2. Remove all of the classifications from the copy other than the ones that are to be edited.
3. Edit the classification to contain the new, updated regular expression.

Currently, classifications can only be edited. However, in a future release the ability to create new classifications will be made available.



Note: Custom classifications take precedence over the default classifications. Similarly, only the first definition found in the classification file will be used.

20 Performance Counter Groups

The CPU contains four performance counters that can be configured to count events inside the CPU. Since there are several dozen individual performance counters, the profiler has grouped similar or complementary performance counters into **Performance Counter Groups**. The profiler allows you to record a particular performance counter group (per-core) each time a sample is taken. This section goes into more detail as which performance counter groups exist, and what kind of data they record.

20.1 Performance Counters Disabled

This is the default selection. When this option is selected, no performance counter data will be recorded with the profile data.

20.2 Per-Core Performance Counter Groups

The following options all contain performance counter settings that are determined exclusively on a per-core basis.

20.2.1 Cycles and Instructions

Records cycles and the number of instructions (categorized between floating-point, integer, and load/store). Also computed are other valuable ratios, such as instructions per cycle (IPC), percent of floating-point instructions, and percent of integer instructions.

Table 20-1: Cycles and Instructions Recorded Data

# Cycles	= PMC1_CYCLE
# Floating-Point Instructions	= PMC3_FPU
# Instructions Completed	= PMC2_INSTRUCTION
# Integer Operations	= PMC4_INTEGER
# Load/Store Instructions	= PMC2_INSTRUCTION - PMC3_FPU - PMC4_INTEGER
% Floating-Point	= 100 * PMC3_FPU / PMC2_INSTRUCTION
% Integer	= 100 * PMC4_INTEGER / PMC2_INSTRUCTION
% Load/Store	= 100 * (PMC2_INSTRUCTION - PMC3_FPU - PMC4_INTEGER) / PMC2_INSTRUCTION
IPC: Instructions per Cycle	= PMC2_INSTRUCTION / PMC1_CYCLE

20.2.2 Branch Prediction Performance

Gives feedback on how well the branch prediction is performing and the breakdown of types of branches. The key piece of data in this group is **% of incorrectly predicted branches**. Note that this group only measures percentages and counts, not cycles.

Table 20-2: Branch Prediction Performance Recorded Data

# All Branches	= PMC2_Bx_FALL_THROUGH + PMC3_Bx_TAKEN
# Correctly Predicted	= PMC1_Bx_UNRESOLVED - PMC4_Bx_MISSED
# Fall-Through Branches	= PMC2_Bx_FALL_THROUGH
# Mispredicted Branches	= PMC4_Bx_MISSED
# Predicted Branches Taken	= PMC1_Bx_UNRESOLVED - PMC2_Bx_FALL_THROUGH
# Unconditional Branches	= PMC2_Bx_FALL_THROUGH + PMC3_Bx_TAKEN - PMC1_Bx_UNRESOLVED
# Unresolved Branches	= PMC1_Bx_UNRESOLVED
% Correctly Predicted	= 100 * (PMC1_Bx_UNRESOLVED - PMC4_Bx_MISSED) / PMC1_Bx_UNRESOLVED
% Mispredicted	= 100 * PMC4_Bx_MISSED / PMC1_Bx_UNRESOLVED

20.2.3 Why Branch Prediction Failed

Measures how many cycles were wasted due to various branch issues, such as waiting for successive branches or waiting on math.

Table 20-3: Why Branch Prediction Failed Recorded Data

# BPU Cycles from Inability to Process New Branches	= PMC4_BPU_THIRD
# BPU Cycles Stalled from LR/CR Dependencies	= PMC3_BPU_LR_CR
# Cycles	= PMC2_CYCLE
# Second Branch Stall Cycles	= PMC1_Bx_STALL_CYCLES
% Stall Branch Dependency	= 100 * PMC3_BPU_LR_CR / PMC2_CYCLES
% Stall Minimum	= 100 * max(PMC1_Bx_STALL_CYCLES, PMC3_BPU_LR_CR, PMC4_BPU_THIRD) / PMC2_CYCLES
% Stall Successive Branches	= 100 * PMC1_Bx_STALL_CYCLES / PMC2_CYCLES
% Stall Too Many Branches	= 100 * PMC4_BPU_THIRD / PMC2_CYCLES

20.2.4 Cache and Memory Performance

Gives feedback on the general performance of the L1 and L2 cache. For more comprehensive L1 measurements that include cycle counts, use the [L1 Cache Performance](#) group.

Table 20-4: Cache and Memory Performance Recorded Data

# L1 D-Cache Misses	= PMC3_DC_MISS
# L1 I-Cache Misses	= PMC2_IC_MISS
# L2 Cache Hits	= PMC1_L2_HIT
# L2 Castouts	= PMC4_L2_CASTOUT
# Total L1 Cache Misses	= PMC2_IC_MISS + PMC3_DC_MISS
# Total L2 Cache Misses	= PMC2_IC_MISS + PMC3_DC_MISS - PMC1_L2_HIT

20.2.5 Is Data Access the Bottleneck?

Helps you determine if data access is the reason a particular function is running slow. If the **Instructions per Cycle (IPC)** metric is below 0.6 and the **% data not found in L2 cache** is more than 30%, then data access is a significant problem. These are somewhat arbitrary cutoffs, so numbers close to these thresholds probably indicate a light or moderate problem.

Table 20-5: Is Data Access the Bottleneck? Recorded Data

# Cycles	= PMC1_CYCLE
# Instructions Completed	= PMC4_INSTRUCTION
# L2 D-Cache Misses	= PMC3_L2_D_MISS
# Load/Store Instructions	= PMC2_LOAD_STORE
% L2 Data Miss	= $100 * \text{PMC2_LOAD_STORE} / \text{PMC4_INSTRUCTION}$
% Load/Store Instructions	= $100 * \text{PMC3_L2_D_MISS} / \text{PMC2_LOAD_STORE}$
IPC: Instructions per Cycle	= $\text{PMC4_INSTRUCTION} / \text{PMC1_CYCLE}$

20.2.6 L1 Cache Performance

Gives feedback on L1 cache performance, breaking it down between data and instruction caches. Interesting metrics in this group include **cycles waiting for memory**, **average cycles waiting for memory**, and **% of time waiting for memory**.

Table 20-6: L1 Cache Performance Recorded Data

# Cycles	= PMC4_CYCLE
----------	--------------

# L1 Cache Load Misses	= PMC1_L1_MISS
# L1 D-Cache Misses	= PMC3_DC_MISS
# L1 I-Cache Misses	= PMC2_IC_MISS
# Memory Wait Cycles	= PMC1_L1_MISS / (PMC2_IC_MISS + PMC3_DC_MISS)
# Total L1 Cache Misses	= PMC2_IC_MISS + PMC3_DC_MISS
% Memory Wait Time	= 100 * PMC1_L1_MISS / PMC4_CYCLE

20.2.7 L2 Cache Performance

Gives feedback on L2 cache performance, breaking it down between data and instruction caches. It also includes the number of writes from L2 to main memory.

Table 20-7: L2 Cache Performance Recorded Data

# L2 Cache Hits	= PMC1_L2_HIT
# L2 Castouts	= PMC4_L2_CASTOUT
# L2 D-Cache Misses	= PMC3_L2_D_MISS
# L2 I-Cache Misses	= PMC2_L2_I_MISS
# Total L2 Cache Misses	= PMC2_L2_I_MISS + PMC3_L2_D_MISS

20.2.8 Cache Misses to Memory

This group helps track down potential slow-downs due to cache lines fetches being serviced by main memory.



Note: PMC1_L1_MISS has a threshold of 63 of set; meaning that it is only incremented if the L1 miss exceeds 63 cycles.

Table 20-8: Cache Misses to Memory Recorded Data

# Cache Misses to Memory	= PMC1_L1_MISS
# Cycles	= PMC4_CYCLE
# Instructions Completed	= PMC2_INSTRUCTION
# L2 D-Cache Misses	= PMC3_L2_D_MISS
% L2 D-Cache Misses to Memory	= 100 * PMC1_L1_MISS / PMC3_L2_D_MISS
% L2 D-Cache Misses to Other Cores	= 100 * (PMC3_L2_D_MISS - PMC1_L1_MISS) / PMC3_L2_D_MISS

IPC: Instructions per Cycle	= $\text{PMC2_INSTRUCTION} / \text{PMC4_CYCLE}$
Misses to Mem per 1000 Instructions	= $1000 * \text{PMC1_L1_MISS} / \text{PMC3_L2_D_MISS}$

20.2.9 Outbound Cache Writes

Focuses on writes from L1 to L2 and writes from L2 to main memory.

Table 20-9: Outbound Cache Writes Recorded Data

# Cycles	= PMC1_CYCLE
# Instructions Completed	= PMC3_INSTRUCTION
# L1 Castouts	= PMC2_L1_CASTOUT
# L2 Castouts	= PMC4_L2_CASTOUT
IPC: Instructions per Cycle	= $\text{PMC3_INSTRUCTION} / \text{PMC1_CYCLE}$

20.2.10 High Cost Cache Writes

Certain cache writes can result in extended write times. These include writes to memory that is marked as "Shared" between core cache lines and conditional writes that fail.

Table 20-10: High Cost Cache Writes Recorded Data

# High Cost L1 Cache Writes	= PMC1_SHARED_STORE
# High Cost L2 Cache Writes	= PMC2_SHARED_STORE
# Retried Conditional Stores	= $\text{PMC4_COND_STORE_INT} - \text{PMC3_COND_STORE}$

20.2.11 Cache Line Push Counts

The Espresso chip allows for sharing of data between each core's individual L1 and L2 caches. This group allows for viewing of the number of times the given core pushed data from its cache to one of the other core's caches.

Table 20-11: Cache Line Push Counts Recorded Data

# L1 Cache Modified Interventions	= $\text{PMC3_L1_MOD_INTERV}$
# L1 Cache Shared Interventions	= $\text{PMC2_L1_SHARED_INTERV}$
# L2 Cache Modified Interventions	= $\text{PMC4_L2_MOD_INTERV}$
# L2 Cache Shared Interventions	= $\text{PMC1_L2_SHARED_INTERV}$
% L1 Interventions	= $100 * (\text{PMC2_L1_SHARED_INTERV} + \text{PMC3_L1_MOD_INTERV}) / (\text{PMC1_L2_SHARED_INTERV})$

	+ PMC2_L1_SHARED_INTERV + PMC3_L1_MOD_INTERV + PMC4_L2_MOD_INTERV)
% L2 Interventions	= 100 * (PMC1_L2_SHARED_INTERV + PMC4_L2_MOD_INTERV) / (PMC1_L2_SHARED_INTERV + PMC2_L1_SHARED_INTERV + PMC3_L1_MOD_INTERV + PMC4_L2_MOD_INTERV)
% Modified Interventions	= 100 * (PMC3_L1_MOD_INTERV + PMC4_L2_MOD_INTERV) / (PMC1_L2_SHARED_INTERV + PMC2_L1_SHARED_INTERV + PMC3_L1_MOD_INTERV + PMC4_L2_MOD_INTERV)
% Shared Interventions	= 100 * (PMC1_L2_SHARED_INTERV + PMC2_L1_SHARED_INTERV) / (PMC1_L2_SHARED_INTERV + PMC2_L1_SHARED_INTERV + PMC3_L1_MOD_INTERV + PMC4_L2_MOD_INTERV)

20.2.12 L1 Sharing Relationships

No explanation available at this time.

Table 20-12: L1 Sharing Relationships Recorded Data

# L1 Cache Modified Interventions	= PMC3_L1_MOD_INTERV
# L1 Cache Shared Interventions	= PMC2_L1_SHARED_INTERV
# L1 Cache Shared Stores	= PMC1_L1_SHARED_STORE

20.2.13 L2 Sharing Relationships

No explanation available at this time.

Table 20-13: L2 Sharing Relationships Recorded Data

# L2 Cache Modified Interventions	= PMC4_L2_MOD_INTERV
# L2 Cache Shared Interventions	= PMC1_L2_SHARED_INTERV
# L2 Cache Shared Stores	= PMC2_L2_SHARED_STORE

20.2.14 TLB Statistics

No explanation available at this time.

Table 20-14: TLB Statistics Recorded Data

# DTLB Misses	= PMC3_DTLB_MISS
# DTLB Search Cycles	= PMC4_DTLB_CYCLE
# ITLB Misses	= PMC2_ITLB_MISS
# ITLB Search Cycles	= PMC1_ITLB_CYCLE
Avg Cycles Per DTLB Miss	= PMC1_ITLB_CYCLE / PMC2_ITLB_MISS
Avg Cycles Per ITLB Miss	= PMC4_DTLB_CYCLE / PMC3_DTLB_MISS

20.3 Chip-Wide Performance Counter Groups

The performance counters in this group make at least partial use of performance counter data that is common for all cores on the Espresso chip. Because of this, much of the data between cores will be almost identical, varying only slightly due to slight differences in sampling time.

20.3.1 Inter-Core Paradox

This option is included for sanity checking. The value returned here should always be 0. If this value is ever non-zero, please save the profile, make a note of the current system settings and what was occurring at the time, and contact support@noa.com as this could indicate a problem.

Table 20-15: Inter-Core Paradox Recorded Data

# CIU Paradox	= PMC1_CIU_PARADOX
---------------	--------------------

20.3.2 Cache Sharing on Chip

No explanation available at this time.

Table 20-16: Cache Sharing On Chip Recorded Data

# CIU Load Requests	= PMC1_CIU_LOAD
# CIU Shared Interventions	= PMC3_CIU_SHARED_INTERV
# CIU Shared Intervention from Two Cores	= PMC4_CIU_SHARED_INTERV
# CIU Store Requests	= PMC2_CIU_STORE

20.3.3 Chip Data Transfers

This option allows for tracking how much data is being moved over the Espresso chip's data bus. It also tracks how many of the transfers are Load operations and how many are Store operations.

Table 20-17: Chip Data Transfers Recorded Data

# BIU Load Requests	= PMC3_BIU_LOAD
---------------------	-----------------

# BIU Store Requests	= PMC4_BIU_STORE
# Cycles	= PMC2_CYCLE
# Data Bus Beats	= PMC1_60X_DBEAT
% BIU Loads	= $100 * \text{PMC3_BIU_LOAD} / (\text{PMC3_BIU_LOAD} + \text{PMC4_BIU_STORE})$
% BIU Stores	= $100 * \text{PMC4_BIU_STORE} / (\text{PMC3_BIU_LOAD} + \text{PMC4_BIU_STORE})$
Bytes Transferred across FSB	= $\text{WidthOfBus} * \text{PMC1_60X_DBEAT}$
MB/s across FSB	= $\text{WidthOfBus} * \text{PMC1_60X_DBEAT} * \text{SpeedOfBus} / \text{PMC2_CYCLE}$

20.4 Correlation Between Sampled Functions and Performance Counters

There is only a very weak correlation between the sampled functions and performance counter data, so be careful not to attribute too much meaning to any coinciding data. For example, if there is a large percentage of data cache write misses at the same time the function `strcmp` was sampled, it is unlikely that the function `strcmp` caused the misses. This is because the performance counter was accumulating data between samples (during which many functions were called) and then at the time of the sample, the program counter happened to be in the `strcmp` function when the performance counter was recorded and cleared. The higher the rate of sampling, the stronger the correlation, but unfortunately the profiler recording code will heavily taint the results (causing extra instruction and data cache misses as well as interfering with the other metrics).

While you might not be able to identify a particular function responsible for a performance counter problem, it might be possible to attribute performance counter data to a particular system or area or the code. For example, if there is a large percentage of data cache read misses right before rendering, you can perhaps attribute the misses to code that executed before rendering, based on knowledge of your own game and the clues given by what functions were called. For example, you might determine that the animation system was responsible for the large percentage of data cache read misses because that particular system executed right before rendering.

With the use of the **Heatmap** mode, the correlation between sampled functions and performance counters is much stronger and might be somewhat more reliable.

20.5 Performance Counter Descriptions

The following table describes the individual CPU performance counters that are used in the above listed groups.



Note: A complete list of all performance counters is documented in the Espresso Developer's User Manual available for download from WarioWorld.

Performance Counter	Explanation
PMC1_60X_DBEAT	The number of data beats on the bus.
PMC1_Bx_STALL_CYCLE	The number of cycles that the dispatcher stalls due to a second branch in the instruction stream.

Performance Counter	Explanation
PMC1_Bx_UNRESOLVED	The number of branches that are unresolved when processed.
PMC1_CIU_LOAD	The number of load requests to the CIU.
PMC1_CIU_PARADOX	The number of times a paradox is detected in the CIU.
PMC1_CYCLE	Processor cycles.
PMC1_ITLB_CYCLE	The number of cycles spent performing table search op. for the ITLB.
PMC1_L1_MISS	The number of loads that miss the L1.
PMC1_L1_SHARED_STORE	The number of times a store to a shared line occurs in the L1 cache.
PMC1_L2_HIT	The number of accesses that hit in the L2 cache.
PMC1_L2_SHARED_INTERV	The number of times the L2 cache supplies shared intervention data.
PMC2_Bx_FALL_TROUGH	The number of fall-through branches.
PMC2_CIU_STORE	The number of store requests to the CIU.
PMC2_CYCLE	Processor cycles.
PMC2_IC_MISS	The number of L1 instruction cache misses.
PMC2_INSTRUCTION	The number of instructions completed.
PMC2_ITLB_MISS	The number of ITLB misses.
PMC2_L1_CASTOUT	The number of L1 castouts to L2.
PMC2_L1_SHARED_INTERV	The number of times the L1 supplies shared intervention data.
PMC2_L2_I_MISS	The number of L2 instruction misses.
PMC2_L2_SHARED_STORE	The number of times a store to a shared line occurs in the L2.
PMC2_LOAD_STORE	The number of completed loads and stores.
PMC3_BIU_LOAD	The number of load requests to the BIU.
PMC3_BPU_LR_CR	The number of cycles the BPU stalls due to LR or CR unresolved dependencies.
PMC3_Bx_TAKEN	The number of predicted branches that were taken.
PMC3_CIU_SHARE_INTERV	The number of times a shared intervention is processed in the CIU.
PMC3_COND_STORE	The number of store conditional instructions completed.
PMC3_DC_MISS	The number of L1 data cache misses.
PMC3_DTLB_MISS	The number of DTLB misses.
PMC3_FPU	The number of instructions completed from the FPU.

Performance Counter	Explanation
PMC3_INSTRUCTION	The number of instructions completed.
PMC3_L1_MOD_INTERV	The number of times the L1 supplies modified intervention data.
PMC3_L2_D_MISS	The number of L2 data misses.
PMC4_BIU_STORE	The number of store requests to the BIU.
PMC4_BPU_THIRD	The number of cycles the BPU cannot process new branches.
PMC4_Bx_MISSED	The number of mispredicted branches.
PMC4_CIU_SHARE_INTERV	The number of times a shared intervention from two cores is processed in the CIU.
PMC4_COND_STORE_INT	The number of store conditional instructions completed with reservation intact.
PMC4_CYCLE	Processor cycles.
PMC4_DTLB_CYCLE	The number of cycles spent performing table searches for DTLB accesses.
PMC4_INSTRUCTION	The number of instructions completed.
PMC4_INTEGER	The number of completed integer operations.
PMC4_L2_CASTOUT	The number of L2 castouts.
PMC4_L2_MOD_INTERV	The number of times the L2 supplies modified intervention data.

21 Command Line Options

In order to provide more control over some of the behavior of the profiler GUI, some command-line options have been provided. The format of the command line should be as follows:

21.1 Command Line Format

```
"Nintendo CPU Profiler.exe" [--options] [NPROF] [ELF] [Executable]
```

[--options]

Zero or more of the options listed in [Available Options](#).

[NPROF]

Enter in a filename of a previously saved NPROF (C:\path\file.nprof) to have it loaded automatically when the profiler starts.

[ELF]

Enter in the path to an ELF file that you would like to use for syncing (C:\path\example.elf). This provides the same functionality as the `--loadelf` option. This tells the GUI that it should perform a **Sync** with the specified file as soon as the **Sync** option becomes available. If Executable is not specified the profiler will attempt to find it automatically.

[Executable]

Enter the path an RPX file to be used for syncing (C:\path\example.rpx). This tells the GUI that it should perform a **Sync** with the specified file as soon as the **Sync** option becomes available. If ELF is not specified the profiler will attempt to find it automatically.



Important: The NPROF parameter should not be specified with the ELF or Executable parameters.

21.2 Available Options

--autosave[=filename]

This tells the GUI that it should save out the results of the profile immediately, rather than requiring a button press. Only the NPROF will be saved if this option is used.

The filename in the command line is optional. If not present, the filename will be based on the currently loaded ELF and will be saved into the same folder as the ELF. The final filenames used will be in the form: `filename-YYYYMMDD-hhmmss.nprof`

--loadelf=filename

(Deprecated) This option has been replaced by just specifying the file directly on the command line. This tells the GUI that it should perform a **Sync** with the specified file as soon as the **Sync** option becomes available.

--saveonerror

(Deprecated) Specifying this option will save an NPROF even if there was a problem in parsing the data. This is primarily intended as a feature to enable sending data to the profiler team in case of a problem in parsing what is believed to be valid data. (It was determined that this feature was too useful to leave as an optional setting. It is now always enabled.)

--script=filename

When the profiler starts up it will automatically run the specified script. This operation will be performed after a specified NPROF is loaded, but may occur before --loadelf has completed.

--shutdown

This option will shutdown the profiler GUI after a profile has been received. It is recommended that this option only be used with the --autosave option.

22 Scripting

The profiler supports loading in Windows PowerShell script files in order to perform automated GUI control. This can be useful if you want to take the same suite of profiles multiple times.

22.1 Windows PowerShell

Windows PowerShell scripting features allow for fairly complex behavior. It is beyond the scope of this document to go into a background of this scripting language. You can find more details on Windows PowerShell scripting from Microsoft.

22.2 Cmdlets

The profiler allows for scripting by creating new cmdlets to be used. All built-in cmdlets should also be available. Extended cmdlets that have been installed separately may not be available.

All of the profiler's cmdlets start with the prefix 'PROFILER-'.



Note: In Windows PowerShell, cmdlets are not case-sensitive.

Profiler cmdlet arguments are space separated and text not enclosed in quotes is converted to upper-case for internal comparisons. This will cause a problem with function names. As such, all function names will need to be enclosed in quotation marks ("""). Quotation marks will also allow for C++ function names with multiple arguments and/or template information to contain spaces if needed.

If the argument name is surrounded by [], passing in the name is optional. If the argument name is omitted than the order of the arguments must be passed in the order shown.

22.2.1 PROFILER-ANALYZE

Asks the profiler to perform an analysis of the profile data.

Story Analysis

This option has the profiler analyze the data using one of the Story Analysis modes discussed in [Story Analysis](#).

```
PROFILER-ANALYZE -Story <StoryType>
```

-Story <StoryType>

One of the types of Story Analysis options. Options listed in [Story Types](#).

Spike Detection

This option has the profiler analyze the data using one the Spike Detection modes discussed in [Spike Detection](#).

```
PROFILER-ANALYZE -SpikeDetection <SpikeDetectionType>
```

-SpikeDetection <SpikeDetectionType>

One of the types of Spike Detection options. Options listed in [Spike Detection Types](#).

Examples

```
PROFILER-ANALYZE -Story Full
PROFILER-ANALYZE -SpikeDetection Rare
```

22.2.2 PROFILER-ASSEMBLY

Controls the options for interacting with the **Assembly** tab.

Set Function

```
PROFILER-ASMFUNCTION -FunctionName <String>
```

-FunctionName <String>

The string provided must match the function name that is displayed in the profiler exactly. If using C++ this will include the arguments and template information.

Show Source

```
PROFILER-ASSEMBLY -ShowSource <Boolean>
```

-ShowSource <Boolean>

Specifies if source code should be shown.

Examples

```
PROFILER-ASSEMBLY -FunctionName "TestFunction"
PROFILER-ASSEMBLY -ShowSource $true
```

22.2.3 PROFILER-CALLSTACKS

Sets whether to record callstacks in a Sampled profile or not.

```
PROFILER-CALLSTACKS [-ShowCallstacks] <Boolean>
```

[-ShowCallstacks] <Boolean>

Specifies if callstacks should be recorded.

Example

```
PROFILER-CALLSTACKS $false
```

22.2.4 PROFILER-COPY

Tells the profiler to copy the contents of a specified tab to the clipboard.

```
PROFILER-COPY [-Tab] <ScriptingTabs> [-ExtendedParam
<ScriptingCopyExtensions>]
```

[-Tab] <ScriptingTabs>

Tab names are listed in [Scripting Tabs](#).

[-ExtendedParam <ScriptingCopyExtensions>]

Optional. Certain tabs have more than one copy option. The extended options are listed in [Copy Extensions](#).

Example

```
PROFILER-COPY Functions
PROFILER-COPY CallTree Selected
PROFILER-COPY CodeCoverage -ExtendedParam Seen
```

22.2.5 PROFILER-FILTER

Sets a filter for the specified tab.

There are a few different forms to this cmdlet. Not all forms are compatible with all tabs.

By Name

This option is the default form of PROFILER-FILTER. It will result in the Filter/Find box on the tab being modified.

```
PROFILER-FILTER [-Tab] <ScriptingTabs> [-Negate] <Boolean> [-RegEx] <Boolean>
[-FilterName] <String>
```

[-Tab] <ScriptingTabs>

Tab names are listed in [Scripting Tabs](#).

[-Negate] <Boolean>

Specifies if the results should be negated.

[-RegEx] <Boolean>

Specifies if the filter is a regular expression.

[-FilterName] <String>

The text to use for the filter. It is recommended that this text be placed in quotation marks, especially if this is a regular expression filter.

Limit

This option will either Limit the results in the Function Tab or allow all results in the Function Tab to be shown.

```
PROFILER-FILTER [-Tab] <ScriptingTabs> -Limit <Boolean>
```

[-Tab] <ScriptingTabs>

Tab names are listed in [Scripting Tabs](#).

Limit <Boolean>

Whether or not to limit the visible results.

Show Selected

This option will set the Function Tab to only show the selected functions.

```
PROFILER-FILTER [-Tab] <ScriptingTabs> -ShowSelected <Boolean>
```

[-Tab] <ScriptingTabs>

Tab names are listed in [Scripting Tabs](#).

ShowSelected <Boolean>

Whether or not to show only the selected items.

By Core

This option will filter the currently visible results by only showing the results from the specified core.

```
PROFILER-FILTER [-Tab] <ScriptingTabs> -Core <Object>
```

[-Tab] <ScriptingTabs>

Tab names are listed in [Scripting Tabs](#).

Core <Object>

Specify either the Integer core number to view, or the string "All" to show all cores.

By Thread

This option will filter the currently visible results by only showing the results from the specified thread.

```
PROFILER-FILTER [-Tab] <ScriptingTabs> -Thread <String>
```

[-Tab] <ScriptingTabs>

Tab names are listed in [Scripting Tabs](#).

Group <String>

The name of the group to filter. Specify "All" to see all groups.

Examples

```
PROFILER-FILTER Functions $false $true "^OS"
PROFILER-FILTER Functions -Negate $false -RegEx $true -FilterName "^OS"
PROFILER-FILTER Functions -ShowSelected $true
PROFILER-FILTER Threads -Core 0
PROFILER-FILTER Functions -Thread "ThreadName"
PROFILER-FILTER Counters -Group "GroupName"
```

22.2.6 PROFILER-INSTRUMENT

Sets the function you wish to instrument based on the specified function name.

To set a function, the symbol must be known by the profiler GUI. This is most easily accomplished by taking a quick Sampled profile before attempting to take an Instrumented profile.

```
PROFILER-INSTRUMENT [-FunctionName] <String>
```

[-FunctionName] <String>

This must match the function name that is displayed in the profiler exactly. If using C++ this will include the arguments and template information.

Example

```
PROFILER-INSTRUMENT "TestFunction"
```

22.2.7 PROFILER-MARKED

Sets whether Marked code should be recorded with the profile data.

Although the GUI allows this value to be set independently on each tab, this cmdlet will set the flag to record Marked code for both Sampled and Instrumented profiles.

```
PROFILER-MARKED [-ShowMarkedCode] <Boolean>
```

[-ShowMarkedCode] <Boolean>

Specifies if marked code should be recorded in the profile data.

Example

```
PROFILER-MARKED $false
```

22.2.8 PROFILER-OPEN

Opens an NPROF file.

```
PROFILER-OPEN [-FileName] <String>
```

[-FileName] <String>

The full path to the NPROF file to open. It is recommended that this path name be enclosed in quotation marks.

Example

```
PROFILER-OPEN "C:/Profiles/SavedProfile1.nprof"
```

22.2.9 PROFILER-PERFCOUNTERS

Sets which performance counters to record with both Sampled and Instrumented profiles.

Although the GUI allows for setting these independently, this cmdlet will set both at the same time.

```
PROFILER-PERFCOUNTERS [-PerfCounters] <ScriptingPerfCounters>
```

[-PerfCounters] <ScriptingPerfCounters>

The Performance Counter values are detailed in [Scripting Perf Counter](#).

Example

```
PROFILER-PERFCOUNTERS Disabled
```

22.2.10 PROFILER-SAMPLE

Sets both the Sample Strategy and Sample Rate for Sampled profiles.

```
PROFILER-SAMPLE [-Strategy] <ScriptingSampleStrategy> [-Rate] <Integer> [-NoWarning]
```

[-Strategy] <ScriptingSampleStrategy>

Strategy values are detailed in [Scripting Sample Strategy](#).

[-Rate] <Integer>

Rate values are detailed in [Rate](#).

[-NoWarning]

This flag keeps a warning from being generated when an invalid Strategy and Rate are combined.

Returns

A Boolean value indicating whether the Strategy and Rate were set correctly.

Example

```
PROFILER-SAMPLE Time 1000
$success = PROFILER-SAMPLE -NoWarning -Strategy Time -Rate 100
```

22.2.11 PROFILER-SAMPLEGRAPH

Controls the Sample Graph options. The primary use-case for this option is for Quick Filters.

ShowCombined

Control whether the combined line should be visible.

```
PROFILER-SAMPLEGRAPH -ShowCombined <Scripting.SampleGraph.CombinedLineDisplay>
```

-ShowCombined <Scripting.SampleGraph.CombinedLineDisplay>

Sets how the combined line is shown. Combined line display options are detailed in [Scripting.SampleGraph.CombinedLineDisplay](#).

SampleDisplay

Control how samples are displayed.

```
PROFILER-SAMPLEGRAPH -SampleDisplay <Scripting.SampleGraph.SampleDisplay>
```

-SampleDisplay <Scripting.SampleGraph.SampleDisplay>

Sets how samples are displayed. Sample display options are detailed in [Scripting.SampleGraph.SampleDisplay](#).

NameDisplay

Control how names are displayed.

```
PROFILER-SAMPLEGRAPH -NameDisplay <Scripting.SampleGraph.FunctionNameDisplay>
```

-NameDisplay <Scripting.SampleGraph.FunctionNameDisplay>

Sets how function names are displayed. Function name display options are detailed in [Scripting.SampleGraph.NameDisplay](#).

Example

```
PROFILER-SAMPLEGRAPH -ShowCombined Only
PROFILER-SAMPLEGRAPH -SampleDisplay Mixed
PROFILER-SAMPLEGRAPH -NameDisplay Short
```

22.2.12 PROFILER-SAVE

Saves the current profile data.

```
PROFILER-SAVE [[-FileName] <String>]
```

[[-FileName] <String>]

Optional. The full path of the file to save. It is recommended that this path name be enclosed in quotation marks. If no value is passed in, the file will be saved to the current directory with the date and time for its name.

Example

```
PROFILER-SAVE
PROFILER-SAVE "C:/Profiles/SavedProfile1.nprof"
```

22.2.13 PROFILER-SELECT

Selects items in the specified tab.

TopGroup

Selects groups of 45 elements at once.

```
PROFILER-SELECT [-Tab] <ScriptingTabs> -Top [[-TopGroup] <Integer>]
```

[-Tab] <ScriptingTabs>

Tab names are listed in [Scripting Tabs](#).

-Top

Specifies that this will be selecting groups of 45 elements at once.

[[-TopGroup] <Integer>]

Which group will be selected. If unspecified, the first 45 elements are selected.

List

Selects elements at the specified indices from the tab (indices run from top to bottom, with the top being 0).

```
PROFILER-SELECT [-Tab] <ScriptingTabs> -List <Integer[]>
```

[-Tab] <ScriptingTabs>

Tab names are listed in [Scripting Tabs](#).

-List <Integer[]>

A list of the indices to select.

DeselectAll

Deselects all currently selected elements.

```
PROFILER-SELECT [-Tab] <ScriptingTabs> -DeselectAll
```

[-Tab] <ScriptingTabs>

Tab names are listed in [Scripting Tabs](#).

-DeselectAll

Specified to deselect all selected elements on the tab.

Example

```
PROFILER-SELECT Functions -Top
PROFILER-SELECT Functions -Top 3
PROFILER-SELECT Functions -List 0 1 4 5
PROFILER-SELECT Functions -DeselectAll

[int[]]items = 0, 3
PROFILER-SELECT Functions -List @items
```

22.2.14 PROFILER-SHOWTAB

Brings a tab into view.

```
PROFILER-SHOWTAB [-Tab] <ScriptingTabs>
```

[-Tab] <ScriptingTabs>

Tab names are listed in [Scripting Tabs](#).

Example

```
PROFILER-SHOWTAB SampleGraph
```

22.2.15 PROFILER-START

Starts taking a profile of the specified type.

```
PROFILER-START [-ProfileType] <ScriptingProfileType>
```

[-ProfileType] <ScriptingProfileType>

The Profile Type values are specified in [Scripting Profile Type](#).

Example

```
PROFILER-START Sampled
PROFILER-START Instrumented
```

22.2.16 PROFILER-STOP

Stops the profile session.

```
PROFILER-STOP
```


Example

```
PROFILER-STOP
```

22.2.17 PROFILER-SYNC

Syncs to the runtime application using the specified ELF.

Rather than use this for automation, it is generally recommended that you use the [Command Line Options](#) to specify an ELF to load automatically.

```
PROFILER-SYNC [[-FileName] <String>]
```

[[-FileName] <String>]

Optional. The full path to the ELF to load. It is recommended that this path name be enclosed in quotation marks. If no ELF is specifically specified, the GUI will attempt to determine what ELF is currently in use to load that.

Returns

A Boolean value if the Sync was successful or not.

Example

```
PROFILER-SYNC
$success = PROFILER-SYNC "C:/Game/bin/NDEBUG/game.elf"
```

22.2.18 PROFILER-TREECONTROL

Tells the profiler to perform actions on the tree present on the specified tab.

```
PROFILER-TREECONTROL [-Tab] <ScriptingTabs> [-Action]
<ScriptingTreeControlActions>
```

[-Tab] <ScriptingTabs>

Tab names are listed in [Scripting Tabs](#).

[-Action] <ScriptingTreeControlActions>

Specifies the action to apply on the tree. Available actions are listed in [Scripting Tree Control Actions](#).

Example

```
PROFILER-TREECONTROL CallTree ExpandAll
PROFILER-TREECONTROL -Tab Info -Action CollapseAll
```

22.2.19 PROFILER-UNITS

Specifies which units should be displayed in the GUI.

```
PROFILER-UNITS [-UnitType] <ScriptingUnitType> [-ShowError] <Boolean>
```

[-UnitType] <ScriptingUnitType>

The Unit Type values are specified in [Scripting Unit Type](#).

[-ShowError] <Boolean>

Whether the error value should be visible.

Example

```
PROFILER-UNITS Percent $false
```

22.2.20 PROFILER-UNSYNC

Disconnects the GUI from the runtime application.

PROFILER-UNSYNC

Example

```
PROFILER-UNSYNC
```

22.2.21 PROFILER-WAIT

Causes the profiler to wait for the specified number of seconds.



Note: This cmdlet is functionally equivalent to the built-in 'SLEEP' command.

PROFILER-WAIT [-Seconds] <Double>

[-Seconds] <Double>

The number of seconds to wait. This value can be fractional.

Example

```
PROFILER-WAIT 0.5
```

22.2.22 PROFILER-WAITFOR

Causes the profiler to wait for a specific event before continuing execution.

PROFILER-WAITFOR [-WaitEvent] <ProfilerWaitEvents>

[-WaitEvent] <ProfilerWaitEvents>

The Event values are specified in [Profiler Wait Events](#).

Example

```
PROFILER-WAITFOR ReadyToProfile
```

22.2.23 PROFILER-CAFE-CORES

Tells the profiler which cores should be recorded in a Sampled profile.

PROFILER-CAFE-CORES <Integer[]>

<Integer[]>

Space separated list of core numbers to record. If a core number is not present in the list, it is disabled for recording. Each core number to record should be specified only once. The order of the numbers does not matter.

Example

```
PROFILER-CAFE-CORES 0 1 2

[int[]]cores = 0, 1
PROFILER-CAFE-CORES @cores
```

22.2.24 PROFILER-CAFE-GPUSTATS

Specifies if GPU Stats should be recorded with the Sampled profile.

PROFILER-CAFE-GPUSTATS [-ShowGpuStats] <Boolean>

[-ShowGpuStats] <Boolean>

Specifies whether GPU statistics should be recorded with the profile data.

Example

```
PROFILER-CAFE-GPUSTATS $true
```

22.2.25 PROFILER-CAFE-OSSTATS

Sets whether to record OS Statistics with a Sampled profile.

PROFILER-CAFE-OSSTATS [-ShowOSStats] <Boolean>

[-ShowOSStats] <Boolean>

Specifies whether OS statistics should be recorded in the profile data.

Example

```
PROFILER-CAFE-OSSTATS $true
```

22.2.26 PROFILER-CAFE-SETDIR

Sets a directory on the Directory Tab to the value specified.

PROFILER-CAFE-SETDIR [-DirType] <ScriptingDirectoryType> [-DirName] <String>

[-DirType] <ScriptingDirectoryType>

The Directory Type values are specified in [Scripting Directory Type](#).

[-DirName] <String>

The full path for that directory. It is recommended that this path name be enclosed in quotation marks.

Example

```
PROFILER-CAFE-SETDIR SLC "C:/cafe_sdk/data/slc"
```

22.3 Argument Values

The values that can be specified in certain cmdlet's arguments are specified below.

22.3.1 Copy Extensions

The Call Tree and Code Coverage tabs have more than one copy functionality.

After each option below, compatible tabs will be listed.

- **All - Tabs:** CallTree, CodeCoverage. Default value. Used to copy all data on the tab.
- **NotSeen - Tabs:** CodeCoverage. Used to copy only the unseen functions list.
- **Seen - Tabs:** CodeCoverage. Used to copy only the seen functions list.
- **Selected - Tabs:** CallTree. Used to copy only the selected functions from the Call Tree.

22.3.2 Profiler Wait Events

Events that can be waited for in the scripting engine.

- CanSync
- ReadyToProfile

22.3.3 Rate

Not all rate values can be specified for all Sample Strategies. To see which rates are compatible with which strategies, please use the GUI. If '0' is specified the value will be set to its default recording value for the strategy chosen.

22.3.4 Scripting Directory Type

Determines which entry on the **Directory Tab** to modify.

- SLC

22.3.5 Scripting Perf Counter

The strings used to denote performance counter groups in a script.

These values correspond with the performance counter groups specified in [Performance Counter Groups](#).

- Disabled
- CyclesInstructions
- BranchPredictionPerf
- WhyBranchPredictionFail

- CacheMemPerf
- IsDataAccessBottleneck
- L1CachePerf
- L2CachePerf
- CacheMissToMem
- OutboundCacheWrites
- HighCostCacheWrites
- CacheLinePushCounts
- L1Sharing
- L2Sharing
- TLBStats
- InterCoreParadox
- CacheSharing
- ChipDataTransfers

22.3.6 Scripting Profile Type

The different types of profiles that can be started.

- `Sampled`
- `Instrumented`

22.3.7 Scripting Sample Strategy

These strings denote sampling strategies in a script.

These values correspond with the strategies presented in [Sampling Strategy and Rate Buttons](#).

- `Time`
- `L1DCacheMiss`
- `L2DCacheMiss`
- `L1ICacheMiss`
- `L2ICacheMiss`
- `L1MissToMem`
- `L1WriteToL2`
- `L2WriteToMem`
- `CyclesStalledL1Miss`
- `BranchMispredict`

22.3.8 Scripting Tabs

A list of the tabs that can be targeted using certain scripting commands.

- `Assembly`
- `CallTree`
- `CodeCoverage`
- `Counters`
- `Functions`
- `Info`
- `Instrumented`
- `Last`
- `Threads`

22.3.9 Scripting Tree Control Actions

A list of the actions that can be performed on trees.

- CollapseAll
- ExpandAll

22.3.10 Scripting Unit Type

The type of units to display in the profiler.

- AvgTime
- Percent
- Samples
- Time

22.3.11 Scripting.SampleGraph.CombinedLineDisplay

Controls how the sample graph shows the combined line.

- Hide
- Show
- Only

22.3.12 Scripting.SampleGraph.NameDisplay

Controls how function and counter names are displayed in the sample graph.

- Hide
- Short
- Long
- Full

22.3.13 Scripting.SampleGraph.SampleDisplay

Controls how Callstack samples are displayed.

- Self
- Total
- Mixed

22.3.14 Spike Detection Types

List of Spike Detection options described in [Spike Detection](#).

- Rare
- Burst
- Flutter
- BadFrame

22.3.15 Story Types

List of story analysis options described in [Story Analysis](#).

- Brief
- Full
- File
- Top

22.4 Other Useful Commands

To maximize the usefulness of the scripting, here are some additional properties, commands, and function calls that may prove useful.

- `[Environment]::CurrentDirectory`

Gets the current directory of the profiler application. This will usually be the folder from which the EXE was launched.

- `[Nintendo.CPU.Profiler.Scripting.ScriptingCore]::NProfFilename`

The entire filename of the currently loaded NPROF file. If no NPROF is loaded, or if this is a new profile that has not yet been saved, this will contain an empty string.

- `[Nintendo.CPU.Profiler.Scripting.ScriptingCore]::ScriptDirectory`

The directory from which the currently executing script was loaded.

- `[Nintendo.CPU.Profiler.Scripting.ScriptingGUI]::ELFDirectory`

The directory of the ELF that is current synced.

23 Wii U CPU Profiler Game API

The Wii U CPU Profiler Game API is required in order to profile a game or application. The API documentation can be found in the supplied [man pages](#).

24 Assembly Files and Functions

For assembly functions that are located in an assembler file (typically a .s file), the .type and .size directives must be used for the profiler to generate correct data. Take the following assembly file:

```
.text

.global ASM_AddNumbers
ASM_AddNumbers:

    add r3, r3, r4
    blr
```

With the ASM_AddNumbers function, the compiler will export the symbol for linking. However, some key information is missing that is required by the profiler in order to ensure proper operation of callstacks. Specifically, the .type and .size directives need to be added in order for the exported symbol to be of type "Function" and the size to indicate the size of the function.

Here is the same assembly file with the directives added:

```
.text

.global ASM_AddNumbers
ASM_AddNumbers:
    .type ASM_AddNumbers, @function

    add r3, r3, r4
    blr
    .size ASM_AddNumbers, $-ASM_AddNumbers
```

25 Troubleshooting

The following is a list of known problems and solutions.

25.1 Profiler crashes

In the case of a crash, a log file is saved out in the same location as the Nintendo CPU Profiler executable. When this happens, please send this log file along with any other important information to Nintendo and we will quickly resolve the issue. Please maintain a version of the application that caused the crash (it may be needed to help diagnose the problem).

25.2 Sample Graph draws very slow and has overlapping frame marker

If you accidentally record heartbeats too fast (faster than every 1ms or so), the **Sample Graph** will draw all of them (causing drawing to become very slow). If this occurs, find the error in your code where you are recording them too often.

25.3 Sample Graph does not appear

If the Sample Graph does not appear, but does display an error message, try following the advice of the error message. Details about DirectX requirements can be found in [Requirements](#).

25.4 Sample Graph appears to be a frame behind

This problem has been reported when either the incorrect graphics drivers are in use, or the graphics drivers are out of date. Please ensure that you have the correct graphics drivers installed for the graphics card and your computer, and that the drivers are up-to-date.


25.5 Sync button stays grayed out

The profiler depends on a sub-application CafeDevKitAvailabilityHost.exe in order to track which dev kits are connected to the computer and which of those the profiler can Sync with. The profiler attempts to launch this sub-application any time it detects that it is not running. To attempt to fix the issue, please try the following (in order):

1. Check for a pop-up window underneath the profiler asking for permission to run an application. (See [Security warning pop-up after starting profiler](#).)
2. Restart the profiler.
3. Manually attempt to start CafeDevKitAvailabilityHost application.

If neither of the above works, please check the profiler 'bin' folder for crash logs for the CafeDevKitAvailabilityHost application. They will be named with the date/time when a crash was detected. Please send any recent logs to your local Nintendo support team so that we can look into the situation further.

25.6 Security warning pop-up after starting profiler



This is related to [Sync button stays grayed out](#). The pop-up window is due to the `CafeDevKitAvailabilityHost` executable being marked as an 'untrusted' application. You can hit the 'Run' button to start it up each time, but as this can be troublesome, use one of the following methods to trust the application.

1. Uncheck the 'Always ask before opening this file' box in the pop-up window.
2. Right-click on the `CafeDevKitAvailabilityHost.exe` file and select 'Properties'. In the 'General' properties window hit the 'Unblock' button.

25.7 Inability to take an Instrumented Profile

If you are seeing the message "Unable to take an instrumented profile as the profiler is not currently operating in Debug Mode" on the Instrumented Tab you will be unable to take a profile. This means that the dev kit has reported to the GUI that your code is marked as execute-only. See [Allowing for Instrumentation](#) for information on how to get Instrumented profiles working.

25.8 Unresolved functions

Due to changes in RPLs in SDK 2.08, these symbols will be more prevalent in the future. See [Unresolved Functions](#) for more information on the exact reasons for unresolved functions. Please also see [Directories Tab](#) for help in setting the correct directories.

25.9 Unable to locate the compiler directory

Certain parts of the profiler make use of the compiler toolset. The profiler checks for the compiler tools by first looking for the environment variable `GHS_ROOT`, and if that cannot be found or the path does not exist, by looking for its install directory as setup by the compiler's installer.

25.10 C++ functions are not demangled

The compiler must be installed to demangle C++ functions. See [Unable to locate the compiler directory](#) for more information.

26 Known Issues

The following are known issues.

26.1 Call Tree cores disappear after Collapse All

In the **Call Tree**, if you expand all, then scroll down a ways, and then collapse all, cores 0 and 1 may disappear from the **Call Tree**. Of those still visible, they will be offset lower than expected with a gap between them and the top of the screen. To resolve the situation, either expand the visible core or use the mouse scroll wheel to attempt to scroll down.

26.2 Call Tree seems wrong

Because of the ABI for Power PC chips, there are currently some cases in which callstacks will come back looking a little off. For reference, the following is the known list of situations where problems can occur:

- The wrong SLC directory is currently set. See [SLC Directory](#) for help on setting the correct directory. We are currently looking for a way to help users identify when this type of problem occurs.
- Sample occurs in a function which cannot be resolved.
- Custom assembly functions which do not conform properly with the ABI. (Example case: The `rand()` function provided with the compiler calls `_savesmall_16` and `_restsmall_16` for its prolog and epilog, respectively.)

26.3 Occasionally wrong SDK function reported with sampling

It is a known bug that particular SDK functions under-report their size in the SDK libraries. This affects the profiler's address-to-function name conversion, sometimes resulting in an SDK function sample having the wrong function name attributed to the sample. It is unknown if or when this may be fixed.

27 Glossary

Buffer

The memory on the dev kit used to store the profile data.

Call Tree

A tree constructed from individual sample callstacks, showing the calling structure of the game during the time it was profiled. This is slightly different from a call graph, since there are no cycles. In a call tree, the same function can appear multiple times, based on where it was called.

Callstack

The trace of function calls discovered at runtime by walking the stack.

Counter

Data from a CPU counter (performance counters) or any other counter described in [Types of Counters](#).

Heartbeat

A periodic event on a particular core that can be manually tracked by using functions built into the platform's API. A heartbeat is similar to a frame marker, but more generic in meaning since it can represent any type of periodic event. Heartbeat markers can be viewed in the **Sample Graph**.

Jitter

Offsets in the sampling times in order to introduce variance in what sections of code are sampled. Without proper jittering there would be a disproportional number of samples in certain sections of the code due to the fixed sampling rates.

Sample by Performance Counter

This is a way to sample functions using a particular performance counter. For example, you could choose to **Sample by L1 Data Cache Miss** at a rate of **Every 500 misses**. For every 500 misses, the profiler will sample what function your game was currently executing.

Sample by Time

This is a way to sample functions based on the specified rate.

Sampling

Periodically stopping the game and recording the current function or callstack.

Sampling Rate

The rate at which samples are taken.

Sampling Strategy

The strategy used to determine when a sample will be taken.

Self

The time spent exclusively in a function, but not any functions it calls.

Sub

The time spent in the functions called by this function ($\text{Sub} = \text{Total} - \text{Self}$).

Sync

Connecting with the dev kit and loading the ELF file. Once the profiler is synced, a profile can be taken.

Total

The time spent in a function and all of the functions it calls.

Units

The units used to show how much time a function has taken. This can be either **Percent**, **Time**, or **Samples**.

Unsync

Disconnecting from the dev kit.

28 Revision History

Version	Revision Date	Description
1.00	2015/07/02	Added System Classification Filter to Functions tab. Changed Reverse Call Tree to Partial Call Tree.
0.19	2015/04/22	Added Last Functions tab. Added Diff feature on Functions tab. Added script commands for new features.
0.18	2014/10/22	Added Checkers. Added Function Details. Added Stack Depth graph. Added Icicle Selected graph. Added more context menus. Added ability to sort on Assembly Tab. Added configurable checkers. Added Text Format Selection to Debug Output. Added Select Function context menu option to Code Coverage. Added script command PROFILER-TREECONTROL. Added details on the Quick Ribbon Bar. Added section on Function Details window operations. Added information on Assembly Tab right-click context menu. Removed option to Run Checkers as this is always performed after a profile is loaded or taken. Renamed 'Debug Output' tab to 'Console' tab. Renamed some options on the Sample Graph right-click context menu.
0.17	2014/04/22	Standardized section layouts. Added Inferred Heartbeats section to Info Tab. Updated Cmdlet list. Updated Glossary. Replaced 1st, 2nd, and 3rd buttons with Select Top button. Fixed some locations where the wrong platform details were referenced. Removed Quick Filter listings of regular expressions. Switched to new format. Added Heartbeats Tab.

Version	Revision Date	Description
		Updated Quick Filter buttons.
		Added Debug Output Tab documentation.
		Added section on Limitations on Selecting Functions to Instrument.
0.16	2013/12/16	Updated requirements. Added troubleshooting for Sync issues.
0.15.1	2013/09/24	Updated Requirements.
0.15	2013/06/19	Added additional profiler workflows. Added Story Analysis section. Added Spike Detection section. Added Sample Graph right click context menu section. Added Mixed description for Sample Graph Self/Total/Mixed drop down section. Added additional GPU Statistics. Added new Directories. Added new Scripting Cmdlets. Added information on Quick Filter customization.
0.14.1	2013/03/13	Added information on --script command line argument.
0.14.0	2013/01/29	Added information on Scripting. Added information on Inferred Heartbeats. Added explanation of buttons in the Quick Filters section of the Home tab of the Ribbon Bar. Added more regular expression examples to Filtering Functions section. Added information on selecting one core's heartbeat to reframe all cores. Added an explanation of the Load per Frame graph in the Sample Graph Tab section. Added information on Sample Gap detection in the Sample Graph. Added more details on GPU Statistics. Removed Graph Drawing Style section (functionality removed).
0.13.0	2012/10/24	Added new Profiler Workflows section. Added new sections for frame selection in the Sample Graph. Added Avg unit button explanation in Home ribbon bar. Added Call Tree right click on node explanation. Added GPU Statistics. Changed Sample Graph right scrollbar explanation to include dragging.

Version	Revision Date	Description
		Updated information on unresolved functions.
0.12.0	2012/08/15	Added new section "User Logged Events" to clarify logging options. Added new heartbeat option VSyncHeartbeat within Sample Graph. Added new known issue.
0.11.0	2012/07/05	Added "Sample by L1 Misses to MEM" explanation to sampling strategy drop down box in Sampled Profile tab. Added "Cache Misses to Memory" performance counter group explanation in the Performance Counter Groups section. Updated Auto Expand explanation in Call Tree tab. Added Combined Graphs section in Sample Graph tab. Added mention that the Call Tree tab can be used to select a function to instrument in the Instrumented Profile tab. Added mention of the vertical scrollbar in the Sample Graph tab. Added mention of directly editing SLC Dir in the Directories tab . Added Known Issue of some SDK function names not being reported correctly due to size misreporting in the SDK.
0.10.0	2012/05/31	Updated Copy behavior on Functions tab and Call Tree tab. Updated number of selected items from 15 to 45. Added multiple select/deselect (Shift-Click, Alt-Click). Added Instrumented graphing explanation on Sample Graph.
0.09.0	2012/04/30	Rewrote the Setup instructions to be more clear and complete. Added ability to Instrument a function (Instrumented Profile tab and Instrumented tab). Added Sample by Performance Counter capability (Sampled Profile tab). Added profile duration counter and timed stop. Added new Usage Statistic. Moved API documentation to separate man pages. Modified Limitations.
0.08.0	2012/03/22	Added Self/Total drop down description to Sample Graph. Added Select Similar button description to Counters tab. Updated "Why Branch Prediction Failed" Perf Counter Group. Added Filter information on new drop down, save and delete buttons. Added mention in Copy Button of Functions tab to explain that the copy now uses tab delimiters and can be pasted into Excel. Added User Logging API. Removed requirement of pmcpu.a library inclusion.

Version	Revision Date	Description
0.07.0	2012/02/15	Added section on Threads tab. Added section on Timeline in Sample Graph. Added mention of Fixed 1ms Intervals in Sample Graph. Added requirement of zlib125.a library inclusion. Added Usage Statistics section. Removed Recorded Sequence tab section. Removed Per-Core Paradox Performance Counter.
0.06.0	2012/01/06	Added section on Assembly tab. Added description of Time/Heartbeat markers in Sample Graph. Added description of Counter Line Style in Sample Graph.
0.05.0	2011/11/15	Added requirement of pmcpu.a library inclusion. Added DirectX requirements. Added section on Sample Graph. Added section on Performance Counter Groups. Added section on Assembly Files and Functions. Added additional information on profiler resource usage.
0.04.0	2011/10/11	Added section on command line arguments and runtime control section in the profiler game API. Added new profiler limitations.
0.03.0	2011/09/13	Updated manual to coincide with the features present in version 0.03 of the Wii U CPU Profiler.
0.01.0	2011/08/26	Initial version.

All company and product names in this document are the trademarks or registered trademarks of their respective companies.

© 2011–2015 Nintendo

The contents of this document cannot be duplicated, copied, reprinted, transferred, distributed, or loaned in whole or in part without the prior approval of Nintendo.